General efficiency-equity trade-offs for dynamic Mirrleesian tax problems*

Charles Brendon
University of Oxford

July 23, 2012

Abstract

This paper shows how to derive a novel set of analytical optimality conditions characterising the solution to an optimal income tax problem in the dynamic Mirrleesian tradition, under the assumption that a 'first-order' approach to incentive compatibility is valid. The method relies on constructing a class of perturbations to the consumption-output allocations of agents in a manner that preserves relevant incentive compatibility constraints. We use it to provide a new decomposition of the efficiency-equity trade-off at the heart of these models, and show how the introduction of type persistence should make the policymaker more tolerant of productive inefficiencies relative to inequity as time progresses. We additionally show that the well-known result that it is optimal to deter savings in this class of model when consumption and labour supply are additively separable in preferences extends to the case in which they are substitutes, but not the more empirically relevant case of complements.

JEL Classification Codes: D82, E61, H21, H24

Keywords: New Dynamic Public Finance, first-order approach, non-separable preferences, inverse Euler condition

*I thank Mikhail Golosov, Charles Gottleib, Albert Marcet, Kevin Roberts, Rick van der Ploeg, Simon Wren-Lewis and, in particular, Antonio Mele for providing helpful comments relating to early versions of this work, as well as seminar audiences at the RES PhD Conference 2012, RES Conference 2012 and CFE Conference 2012. All errors are mine.

1 Introduction

This paper presents a novel approach to analysing dynamic optimal income tax models, and explores the specific insights that this method yields for optimal public policy. The setup that we adopt is in the New Dynamic Public Finance tradition. That is, it builds on the seminal contributions of Mirrlees (1971) and Diamond and Mirrlees (1978) to consider the best taxation strategy over time for a utilitarian government that is able to observe its citizens' income and saving levels, but not their underlying abilities. This is a problem that has been widely studied by macroeconomists in recent years, with important recent insights provided by Golosov, Troshkin and Tsyvinski (2011) and Farhi and Werning (2011).[1]

Our innovation is to provide a new, intuitive analytical characterisation of the full set of necessary optimality conditions for these problems. Working under the assumption that the 'first-order approach' is valid – so that the set of incentive compatibility constraints that binds at the optimum is known – we construct a class of perturbations to the optimal scheme that will remain within the (binding) constraint set. The incentive-feasibility of these perturbations implies they cannot raise a resource surplus, and optimality conditions then follow by conventional logic. These conditions are far simpler in form than alternatives that exist in the literature, and have natural interpretations in terms of the e'fficiency-equity trade-off' that is so often observed in models of optimal taxation. Notably, they also provide a substantial simplification of the canonical optimality conditions for static Mirrleesian income tax problems, relative to those derived by Mirrlees himself and by Saez (2001).

One way to understand this technical contribution is in showing how a satisfactory version of the calculus of variations can be devised for dynamic models of asymmetric information – contrasting with the recursive dynamic programming approaches more common in the literature. Similar logic has already been applied by Golosov, Kocherlakota and Tsyvinski (2003) and Kocherlakota (2005), among others, to obtain the 'inverse Euler condition' optimality requirement in isolation;[2] the current paper shows how this logic can be extended much more broadly.

The main limitation of our method is its dependence on the first-order approach to incentive compatibility – whereby local incentive constraints are substituted for global. At present the only known conditions under which the validity of this approach

[1] A thorough survey of the literature can be found in Kocherlakota (2011).

[2] This condition is a necessary requirement for dynamic optimality when consumption and labour supply are additively separable in preferences. It links the inverse of the marginal utility of consumption in a given period to its expected value in the next.

can be assured impose requirements on the properties of solutions obtained under it;[3] these conditions cannot, therefore, be confirmed *ex-ante*, based on the model's priors alone. But a number of recent papers have made good progress under the assumption of first-order incentive compatibility – using this to simplify a recursive programming problem, and confirming its validity *ex-post* in any numerical examples. These include Kapička (2011a), Broer, Kapička and Klein (2011), Golosov, Troshkin and Tsyvinski (2011) and Farhi and Werning (2011). We follow them, and thus retain their caveats.

In terms of policy, some of our most interesting results relate to optimal savings wedges. By generalising the inverse Euler condition to situations in which there is non-separability in preferences between consumption and labour supply we are able to show that optimal taxes should always deter savings, in a sense that we make clear, when consumption and labour supply are either substitutes or separable, but that this does not generalise to the empirically-relevant case in which they are complements.[4] In a related result, we show that the long-run 'immiseration' that can characterise dynamic Mirrlees economies with infinitely-lived dynasties and preference separability again generalises to the case of substitutes, but not to that of complements.

Our work in this regard develops that of Golosov, Troshkin and Tsyvinski (2011), who have shown how to obtain an alternative generalisation of the inverse Euler condition in the non-separable case, linking an agent's marginal utility of consumption to the future derivative of an *indirect* utility function that they define. Our analysis suggests that still greater clarity can be obtained if *output* is perturbed in a specific way alongside any given consumption change. In particular, this allows the model's key dynamic trade-off to be expressed through the arguments of the direct utility function alone, from which our results on savings wedges and immiseration follow.

In addition to savings distortions, another important economic variable that features in dynamic optimal tax analyses is the degree of productive distortion – that is, the 'labour wedge', between an agent's within-period marginal rate of substitution between consumption and production, and the marginal rate of transformation. The recent paper by Farhi and Werning (2011) places particular focus on this object, characterising its dynamics when type processes are persistent, and focusing particular attention on the case of separable preferences when labour is supplied isoelastically.

[3] See, for instance, Theorem 5 in Kapička (2011a). The results and discussion in Pavan, Segal and Toikka (2011) provide a more general exploration of the first-order approach across a broad class of dynamic models.

[4] Basu and Kimball (2002), for instance, emphasise that consumption and labour supply must be complements if a low elasticity of intertemporal substitution is to be reconciled with no net effects on labour supply of long-run growth in the real wage.

In this paper we provide an analogous dynamic relationship to theirs that holds under general preference specifications, and show how it can be interpreted in terms of a dynamic efficiency-equity trade-off.[5] In particular, we show that when types are persistent the policymaker should be willing to accept more productive inefficiencies through time relative to a particular measure of the degree of inequality.

Moreover, we are able to be much more precise than other authors about the nature of the extra productive inefficiencies that the policymaker should be willing to accept. By focusing on optimality conditions that must hold across types within any given period, we show that there is greatest incentive to accept extra productive inefficiency at points where the conditional (Markovian) type distribution is most sensitive to variations in past type.

This incentive to induce additional inefficiencies derives entirely from type persistence: by contrast, when skills are iid the set of 'intratemporal' optimality conditions characterising allocations is identical to the set of conditions that must hold in a static optimal income tax model, providing an important mapping between the traditional 'public economics' and more recent 'mechanism design' literatures.[6]

Our exploration of this area contrasts slightly with the study of the labour wedge contained in Golosov, Troshkin and Tsyvinski (2011). These authors derive an expression for the wedge by use of the maximum principle, in a manner deliberately analogous to that of Mirrlees (1971) and Saez (2001). Golosov, Troshkin and Tsyvinski emphasise the additional complexity added to the Mirrleesian optimality requirements when the problem is cast in a dynamic setting. The point of our paper is to show that this extra complexity is limited in its scope. At most the dynamic setting breaks one of the optimality conditions from the static problem. All other differences follow from the dynamic casting of the model's *constraints*, and need not be direct concerns when assessing optimality.

2 Model setup

The economy is populated by a large number of infinitely-lived agents (dynasties), each indexed by a position on the unit interval. Labour is the only factor of production and there are no firms – so agents can be thought of as directly choosing the level of output that they produce each period via their labour supply decision. Their

[5] To be clear, Farhi and Werning provide a general expression for the evolution of this wedge when preferences are non-separable, but only as a function of co-state variables. Our result relates directly to the arguments of the direct utility function, and thus is comparatively easy to interpret.

[6] The distinction is drawn by Diamond and Saez (2011).

preferences over output and consumption profiles from time t onwards are described by U_t:

$$U_t = E_t \sum_{s=0}^{\infty} \beta^s u\left(c_{t+s}, y_{t+s}; \theta_{t+s}\right) \tag{1}$$

where $u : \mathbb{R}^3_+ \to \mathbb{R}$. c_{t+s} and y_{t+s} are, respectively, the agent's consumption and output levels in period $t+s$, $\beta \in (0,1)$ is the dynasty's time preference parameter, and θ_{t+s} is an idiosyncratic productivity parameter. The productivity parameter belongs to a set $\Theta \subset \mathbb{R}$, which is time- and history-invariant. For the entirety of this paper we work under the assumption that Θ is a countable set with cardinality N,[7] which is the simplest setting in which to present the main arguments; generalising to the case in which Θ is a continuous interval of $\mathbb{R}$ is non-trivial, but can be done. Expectations are taken across all stochastic variables relevant to the equilibrium evolution of the agent's utility.

We analyse the model as if nature is responsible at the start of time for drawing a distinct element for each dynasty from the infinite-dimensional product space Θ^∞, according to some probability measure on Θ^∞, π_Θ. These draws are iid across dynasties. At the start of each period agents are informed of their within-period productivity, so that at time t they are aware of their complete history of draws to date, $\theta^t \in \Theta^t$, where θ^t is a t-length truncation of θ^∞. This information is private knowledge to the agent, so policymakers must provide sufficient incentives to prevent mimicking across types in any tax system that is established.

2.1 Restrictions on utility

We assume that the utility function is twice continuously differentiable in all of its arguments, with $u_c > 0$, $u_y < 0$, and $u_\theta > 0$, and that the partial Hessian $\begin{bmatrix} u_{cc} & u_{cy} \\ u_{cy} & u_{yy} \end{bmatrix}$ is negative definite for any given θ. Consumption and leisure are both assumed to be normal goods, and we additionally impose Inada conditions: $\lim_{c\to\infty} u_c\left(c, y; \theta\right) = 0$ and $\lim_{c\to 0} u_c\left(c, y; \theta\right) = \infty$ for all non-zero, finite (y, θ) pairs, and $\lim_{y\to\infty} u_y\left(c, y; \theta\right) = -\infty$ and $\lim_{y\to 0} u_y\left(c, y; \theta\right) = 0$ for all non-zero, finite (c, θ) pairs. These conditions ensure an interior solution obtains at all finite horizons.

To maintain the interpretation of θ as an unobservable productivity level we impose a number of additional restrictions on utility. First, if consumption and utility are jointly held constant as θ is changed then labour supply must implicitly also be

[7] We allow the possiblity that $N = \infty$.

being held fixed – and thus the marginal utility of consumption should likewise be constant. This implies:

$$u_{c\theta} - u_{cy}\frac{u_\theta}{u_y} = 0 \tag{2}$$

Similarly, marginal changes to θ should reduce the marginal disutility associated with a unit of extra output when consumption and utility (and thus, implicitly, labour) are jointly held constant. This can be shown to imply:

$$u_{y\theta} > u_{yy}\frac{u_\theta}{u_y} > 0 \tag{3}$$

A variant on the Spence-Mirrlees single-crossing condition is additionally imposed; this helps to justify the first-order approach in what follows:

$$u\left(c'', y''; \theta''\right) - u\left(c', y'; \theta''\right) > u\left(c'', y''; \theta'\right) - u\left(c', y'; \theta'\right) \tag{4}$$

whenever $c'' > c'$, $y'' > y'$ and $\theta'' > \theta'$. It is easy to show that this implies (but is not limited to) the usual diminishing marginal rate of substitution between consumption and labour supply:

$$-\frac{u_y\left(c, y; \theta''\right)}{u_c\left(c, y; \theta''\right)} < -\frac{u_y\left(c, y; \theta'\right)}{u_c\left(c, y; \theta'\right)} \tag{5}$$

for all (c, y) pairs and any $(\theta'', \theta') \in \Theta^2$ such that $\theta'' > \theta'$.

2.2 Policy

The policymaker is assumed to observe output and savings, but not underlying types. Since consumption can be inferred from output and saving levels we can treat consumption and output as the main observables without loss of generality. The role of policy is to choose, at the start of time, effective tax schedules for all future periods that will link an individual's consumption to their output, conditional on their history of actions to date. The purpose of this choice is to maximise a social welfare function, defined across the set of possible equilibrium allocations.

Since the revelation principle will apply in this setting,[8] we may restrict policy choice to direct revelation mechanisms that deliver consumption and output bundles to individuals as functions of direct current and past productivity reports. We denote by $\sigma_t^i : \Theta^t \to \Theta$ individual i's report at time t as a function of their actual productivity

[8] We seek a Bayes-Nash equilibrium of the game played between the policymaker and all individuals whose types may be drawn from Θ^∞. The revelation principle states that any such equilibrium can be supported by a direct revelation mechanism.

(where this function is measurable with respect to θ^t), by $\sigma^{i,t}:\Theta^t\rightarrow\Theta^t$ the sequence of such reporting plans up to time t, and by $\sigma^i:\Theta^\infty\rightarrow\Theta^\infty$ a complete dynamic reporting plan. We occasionally refer to $\sigma^{i,t}(\cdot)$ as the t-truncation of $\sigma^i(\cdot)$.[9]

For the remainder of the paper we follow the majority of the literature and assume a utilitarian policy criterion, assessed from the perspective of the initial time period. The policymaker's primal choice problem is:

$$\max_{\{c_t(\theta^\infty),y_t(\theta^\infty)\}_{t=1}^{\infty}}\int_{\Theta^\infty}\sum_{t=1}^{\infty}\beta^{t-1}u\left(c_t\left(\theta^\infty\right),y_t\left(\theta^\infty\right);\theta_t\right)d\pi_\Theta\left(\theta^\infty\right) \tag{6}$$

subject to $c_t(\theta^\infty)$ and $y_t(\theta^\infty)$ being measurable with respect to θ^t, together with the incentive compatibility constraints:

$$\begin{aligned}&\int_{\Theta^\infty}\sum_{s=0}^{\infty}\beta^s u\left(c_{t+s}\left(\theta^\infty\right),y_{t+s}\left(\theta^\infty\right);\theta_{t+s}\right)d\pi_\Theta\left(\theta^\infty|\theta^t\right)\\ &\geq\int_{\Theta^\infty}\sum_{s=0}^{\infty}\beta^s u\left(c_{t+s}\left(\sigma\left(\theta^\infty\right)\right),y_{t+s}\left(\sigma\left(\theta^\infty\right)\right);\theta_{t+s}\right)d\pi_\Theta\left(\theta^\infty|\theta^t\right)\end{aligned} \tag{7}$$

which must hold for all t, all θ^t, and all functions $\sigma:\Theta^\infty\rightarrow\Theta^\infty$ whose s-truncations $\sigma^s(\cdot)$ are measurable with respect to θ^s for all $s\geq 1$; and finally the resource constraint:

$$\int_{\Theta^\infty}\left[c_t\left(\theta^\infty\right)-y_t\left(\theta^\infty\right)\right]d\pi_\Theta\left(\theta^\infty\right)+A_{t+1}=R_tA_t \tag{8}$$

where A_t is the quantity of real bonds that the policymaker purchases for time t, each paying R_t units of real income in that period. The initial value R_1A_1 is given. Dynamic solvency requires that $\lim_{s\rightarrow\infty}\left[\left(\prod_{r=1}^{s}R_{t+r}^{-1}\right)A_{t+s}\right]=0$.

3 The first-order approach to incentive compatibility

The set of constraints implied by condition (7) is extremely large. For all type histories at each point in time it requires that truthful reporting should be superior to mimicking each of the $N-1$ other types in Θ. But in general only a very small subset of these constraints will be binding at the optimum, and it will help the analysis substantially if we can suppress the remainder. This is what is done under the

[9] We economise on functional arguments and scripts wherever this will not cause confusion.

'first-order approach' to incentive compatibility, increasingly widely applied in the dynamic Mirrleesian literature.

For all periods $t \geq 1$, define $\sigma_{m,t} : \Theta^t \to \Theta$ as the reporting strategy associated with mimicking a type one lower than the truth in period t:

$$\sigma_{m,t}\left(\theta^{t-1}, \theta_t\right) = \theta'_t$$

where $\theta'_t = \max\{\theta \in \Theta : \theta < \theta_t\}$. If $\theta_t = \min\{\theta \in \Theta\}$ then we normalise $\sigma_{m,t}\left(\theta^{t-1}, \theta_t\right) = \theta_t$. Let $\sigma_{m(t)} : \Theta^\infty \to \Theta^\infty$ be the reporting strategy given by truthful reporting in all periods apart from t, when $\sigma_{m,t}\left(\theta_t\right)$ is applied.[10]

We further define the value function $W\left(\sigma_t\left(\theta^t\right); \theta_t, \sigma^{t-1}\left(\theta^{t-1}\right)\right)$, with $W : \Theta \times \Theta \times \Theta^{t-1} \to \mathbb{R}$ specifying the maximum lifetime utility that could be expected for an agent whose past reports were given by $\sigma^{t-1}\left(\theta^{t-1}\right)$, whose current productivity is θ_t and whose current report is $\sigma_t\left(\theta^t\right)$. If global incentive compatibility holds then for a given $\left(\theta^t, \sigma^{t-1}\right)$ pair this function must attain a maximum where $\sigma_t\left(\theta^t\right) = \theta_t$. Thus instead of choosing directly from among the (difficult to characterise) set of allocations for which (7) is explicitly asserted for all admissible functions σ, it may be enough instead to impose a more limited restriction that *local* movements in the value of $\sigma_t\left(\theta^t\right)$ away from θ_t do not raise utility, conditional on any past report history.

In continuous-type models this would generally amount to imposing a zero restriction on the derivative of W with respect to its first argument. Here types are discrete, so in principle we could apply constraints ruling out local minicking both 'upwards' – i.e., of those who are marginally *more* productive – and 'downwards'. But it is well-known that the first-best allocation in these models (when the policymaker is able to observe types perfectly) involves utility *decreasing* in type.[11] Thus we conjecture that 'upwards' constraints will not bind, and simply impose:

$$W\left(\theta_t; \theta_t, \sigma^{t-1}\right) \geq W\left(\sigma_{m,t}\left(\theta_t\right); \theta_t, \sigma^{t-1}\right) \tag{9}$$

This is clearly necessary but not sufficient for full incentive compatibility. But in some circumstances sufficiency *will* be guaranteed. We prove the following in the appendix:

Proposition 1 ***Sufficiency of first-order approach:*** *Suppose that under a given*

[10] Under Markovian productivities the optimality of a current report does not depend on the veracity of past reports. Thus we do not need to give special consideration to the question of whether a deviation at t will change the optimality of reporting truthfully at $t+1$ or subsequently.

[11] Dasgupta (1982) provides a useful discussion of the intuition behind this result.

policy strategy the value function $W\left(\sigma_t\left(\theta^t\right);\theta_t,\sigma^{t-1}\left(\theta^{t-1}\right)\right)$ satisfies increasing differences in $\left(\sigma_t\left(\theta^t\right),\theta_t\right)$, so that for given $\sigma^{t-1}\left(\theta^{t-1}\right)$ the inequality

$$W\left(\widehat{\theta}_t'';\theta_t'',\sigma^{t-1}\right)-W\left(\widehat{\theta}_t';\theta_t'',\sigma^{t-1}\right)>W\left(\widehat{\theta}_t'';\theta_t',\sigma^{t-1}\right)-W\left(\widehat{\theta}_t';\theta_t',\sigma^{t-1}\right)$$

holds for all $\left(\widehat{\theta}_t'',\widehat{\theta}_t',\theta_t'',\theta_t'\right)\in\Theta^4$ such that $\widehat{\theta}_t''>\widehat{\theta}_t'$ and $\theta_t''>\theta_t'$. Then if condition (9) holds with equality for all $\theta_t\in\Theta$, we must have $W\left(\theta_t;\theta_t,\sigma^{t-1}\right)>W\left(\theta_t'';\theta_t,\sigma^{t-1}\right)$ for all $\theta_t''\in\Theta\backslash\left(\sigma_{m,t}\left(\theta^t\right),\theta_t\right)$.

This is a natural translation to our discrete-type setting of Theorem 5 in Kapička (2011a). Like that result it should really be seen as an intermediate step towards full sufficiency conditions for the first-order approach, since the value function in any given setting will itself depend endogenously upon the chosen policy. But it at least provides a way to check the applicability of the arguments that follow. Moreover, combined with the single-crossing condition we have enough here to assert something much stronger about the iid case:

Corollary 1 *Suppose that agent-level productivities follow an iid process, and that the single-crossing condition (4) applies. Then provided a given policy strategy requires higher-type agents with a given history to produce higher output quantities than lower-type agents with the same history, and simultaneously provides them with higher consumption, condition (9) holding with equality is sufficient for incentive compatability.*

Proof. When productivity shocks are iid, agents' values from $t+1$ on for a given report σ^t are identical in expectation at t, irrespective of their true types. Hence increasing differences will follow provided we have:

$$u\left(\widehat{\theta}_t'';\theta_t'',\sigma^{t-1}\right)-u\left(\widehat{\theta}_t';\theta_t'',\sigma^{t-1}\right)>u\left(\widehat{\theta}_t'';\theta_t',\sigma^{t-1}\right)-u\left(\widehat{\theta}_t';\theta_t',\sigma^{t-1}\right)$$

for all $\left(\widehat{\theta}_t'',\widehat{\theta}_t',\theta_t'',\theta_t'\right)\in\Theta^4$ such that $\widehat{\theta}_t''>\widehat{\theta}_t'$ and $\theta_t''>\theta_t'$. The result is then a direct implication of the single crossing condition, given that output and consumption are increasing in type. ■

Whilst this result still depends on output and consumption increasing in type, this is a much more straightforward condition to confirm *ex-post* than an increasing differences restriction on the value function.

In what follows we refer to the problem of policy choice from among the set of direct revelation mechanisms satisfying condition (9) as the 'relaxed' problem, in contrast with the 'general' problem that imposes (7) directly for all $(\theta_t, \sigma^{t-1}) \in \Theta^t$. Our focus will be on the properties of the solution to this relaxed problem. We therefore make the following assumption throughout:

Assumption: *The solution to the relaxed problem also solves the general problem.*

4 A perturbation approach

We now build up some of the formal apparatus required for our analytical approach.

Conditional upon a particular history of reports σ^{t-1}, an agent's period-t consumption and output allocations under any optimal direct revelation mechanism can be described by an $N \times 2$ matrix $X_t(\sigma^{t-1})$, with each row in this matrix corresponding to a given period-t report,[12] and the columns listing, in turn, associated consumption and output levels. Our main strategy is to perturb these allocations by the addition of one or more of a particular set of $N \times 2$ matrix-valued functions, which in the generic case we denote by $\Delta(\delta)$ (with $\Delta : \mathbb{R} \to \mathbb{R}^{2N}$) for some relevant parameter δ. The perturbed allocations are then described by $X_t(\sigma^{t-1}) + \Delta(\delta)$. The Δ functions are always normalised such that $\Delta(0) = 0$.

In certain cases we will additionally allow changes to be spread through time, with the consumption and output of agents with a common reported history θ^{t-1} changed at $t-1$ as well as at t, according to an analogous function $\Delta_{-1}(\delta)$ (with $\Delta_{-1} : \mathbb{R} \to \mathbb{R}^2$).

We wish to construct these Δ and Δ_{-1} functions so that they satisfy the following three properties:

1. The perturbed allocations remain candidate solutions to the relaxed problem.
2. $\Delta(\delta)$ and $\Delta_{-1}(\delta)$ should be both continuous and continuously differentiable in an open neighbourhood of $\delta = 0$.
3. Expected lifetime utility averaged across all agents should remain constant from the perspective of the initial period for all δ in the neighbourhood of $\delta = 0$.

Under the first-order approach the first property requires that any increase in the utility that an agent of type θ_t^n would obtain from mimicking an agent of type θ_t^{n-1} is offset by at least an equal increase in the utility that the agent of type θ_t^n receives

[12] We assume that these are ordered in ascending values for θ, with the lowest (reported) period-t type's allocation in the first row of X and the highest type's in the Nth row.

from truthful reporting.[13] This ensures that if the original allocation satisfied the constraint set of the relaxed problem then the perturbed allocation must likewise.

The second condition is needed for the perturbations to be applied symmetrically as δ moves above and below 0. It provides a substantial obstacle relative to the first: if we know that incentive compatibility constraints bind downwards then we know it always going to be possible to *increase* the utility of the highest type alone, or of the top n types in sufficiently skewed proportions, so that these constraints will remain strictly satisfied. This could be done, for instance, by the provision of extra consumption to higher-type agents. But perturbations of this form will only ever give us *in*equality restrictions – to the effect that the net marginal cost of changing outcomes in such a manner must be weakly positive. Unless a symmetric downward shift in the utility of high types is possible, this cannot be converted into a first-order condition that is stated with *equality*.

Under the third condition we assume allocations are changed such that the average value across agents of expected lifetime utility remains constant from the perspective of the very first period. To a utilitarian policymaker, assessing outcomes from the perspective of the initial period, there will thus be no direct loss or gain from the perturbation. A necessary condition for the original allocations to have been optimal is then that the marginal effect on the resource cost of utility provision associated with any admissible perturbation should be zero.

4.1 Example: changes at the top

A simple example of a perturbation that satisfies all three of the above requirements is a change in the allocation of the 'top' agent for any given reporting history, such that this agent remains at a constant utility level. That this is incentive-feasible under the relaxed problem derives from the fact that no other agent envies the allocation of the highest type in equilibrium. It provides a straightforward introduction to our approach, and re-states in the dynamic case the well-known 'No distortion at the top' maxim of Mirrlees (1971).

Proposition 2 ***No distortion at the top:*** *In all time periods $t \geq 1$ and for all past reporting histories θ^{t-1}, the constrained-optimal allocation (c_t, y_t) for the agent who reports $\theta_t' = \max\{\theta \in \Theta\}$ satisfies $u_c(c_t, y_t; \theta_t') = -u_y(c_t, y_t; \theta_t')$.*

[13] We use superscripts here to index the agents' types within the set Θ, with θ_t^n increasing in $n \in \{1, ..., N\}$

Proof. Consider a perturbation to the allocation $X_t(\theta^{t-1})$ given by the $N\times 2$ matrix of functions $\Delta:\mathbb{R}\to\mathbb{R}^{2N}$ such that the nth row of $\Delta(\delta)$ equals $(0,0)$ for all $n\in\{1,...,N-1\}$ and the Nth row equals $(\delta, f(\delta))$, with the function $f:\mathbb{R}\to\mathbb{R}$ defined implicitly by:

$$u(c_t+\delta, y_t+f(\delta);\theta_t') = u(c_t, y_t;\theta_t') \tag{10}$$

By construction this change keeps constant the utility of all truth-telling agents in all time periods. The utility of agents who report θ_t' when not of type θ_t' may change, but these objects do not feature in the relaxed constraint set. Hence we remain within that constraint set, and so by assumption cannot improve upon the optimal outcome. Since utility is held constant for all agents the value of the policymaker's objective remains unchanged as δ is varied away from $\delta=0$. The net impact on resources available in period t is $\pi_\Theta(\theta_t'|\theta^{t-1})\pi_\Theta(\theta^{t-1})[f(\delta)-\delta]$. If the original allocation is optimal then the *marginal* impact on resources as δ moves away from zero must be zero. Hence we have:

$$f'(0) = 1 \tag{11}$$

Since utility for a highest-type truth-teller is unchanged by the perturbation we can assert the total derivative:

$$u_c(c_t, y_t;\theta_t') + u_y(c_t, y_t;\theta_t') f'(0) = 0 \tag{12}$$

The result follows immediately. ■

5 Optimal savings distortions

Whilst this 'no distortion at the top' result is interesting, it is not all that surprising given that the top agent is never envied in the relaxed problem. It also clearly relies on the assumption of a finite upper support on the type distribution.[14] More substantive insights come from considering the model's optimality requirements elsewhere, as we do in the remainder of the paper.

We first turn to dynamic optimality, and in particular to optimal savings wedges. When separability holds between consumption and labour supply we have the well-known 'inverse Euler condition', which can easily be used to show that it is optimal to deter savings at the margin – so that the current marginal utility of consumption is

[14] Diamond (1998) and Saez (2001) have shown that optimal tax rates in the static Mirrlees problem will converge to a strictly positive constant at the upper end of an income distribution with sufficiently fat tails and no upper support, including a Pareto distribution.

below its expected future value (after allowing for interest and discounting). But the implications of non-separability are still relatively unclear.[15] This sub-section shows how to arrive at a relatively simple generalisation of the inverse Euler result, and explores its implications for optimal savings wedges.

5.1 Definition of α function

To aid the subsequent presentation we define the function $\alpha(c, y; \theta)$, with $\alpha : \mathbb{R}_+^2 \times \Theta \to \mathbb{R}$, as follows:

$$\alpha(c, y; \theta) := \frac{u_c(c, y; \theta) - u_c(c, y; \theta')}{u_y(c, y; \theta') - u_y(c, y; \theta)} \tag{13}$$

provided $\theta \neq \max\{\theta \in \Theta\}$, where $\theta' = \min\{\theta'' \in \Theta : \theta'' > \theta\}$. If $\theta = \max\{\theta \in \Theta\}$ we simply let $\alpha(c, y; \theta) = 0$.

This α function is useful in understanding the arguments that follow. It gives the marginal increase in output per unit marginal increase in consumption if the combined perturbation is to have an equal impact on utility for agents of both types θ and θ' at the given allocation. Thus it shows how to provide utility at the margin along a dimension in consumption-output space that will ensure both truth-tellers (θ-types) and would-be mimickers (θ'-types) receive the *same* utility increment.

Note that the denominator in (13) is strictly positive. If consumption is additively separable in utility then the numerator is zero, so $\alpha = 0$ always holds. When consumption and labour supply are Edgeworth complements we will have $\alpha > 0$.[16] Intuitively this is because under complementarity the (truth-telling) lower-type agents will receive a greater marginal benefit from a unit increase in consumption at any given allocation than the (mimicking) higher-type agents – because of the higher number of hours the lower types are working to produce the given output level. To offset this disparity one must exploit the higher marginal *dis*utility of additional output for lower types, by requiring that greater production should accompany the increased consumption. Conversely, when consumption and labour supply are Edgeworth substitutes we must have $\alpha < 0$.

[15] The work of Golosov, Troshkin and Tsyvinski (2011) has provided insight into the problem, but these authors state optimality conditions in terms of an indirect utility function in the non-separable case. Here we provide the natural analogue in terms of direct utility functions, which proves more tractable in making qualitative statements about optimal distortions.

[16] Formally, we take consumption and labour supply to be Edgeworth complements if and only if $u_{cy} > 0$, and Edgeworth substitutes if and only if $u_{cy} < 0$. Since these cross-partials hold θ fixed, higher output is equivalent to higher labour supply. Note that equation (2) further implies $u_{c\theta} < 0$ for Edgeworth complements and $u_{c\theta} > 0$ for Edgeworth substitutes.

5.2 A generalised inverse Euler condition

Using this α function we are able to state the following:

Proposition 3 ***Generalised inverse Euler condition:*** *For all time periods $t \geq 1$ and for any reporting history $\theta^t \in \Theta^t$, the constrained-optimal allocations $\left(c_t\left(\theta^t\right), y_t\left(\theta^t\right)\right)$ and $X_{t+1}\left(\theta^t\right)$ for agents who report θ^t satisfy the following condition:*

$$R_{t+1}\beta\frac{1-\alpha\left(\theta_t\right)}{u_c\left(\theta_t\right)+u_y\left(\theta_t\right)\alpha\left(\theta_t\right)} = \sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta^t\right)\frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)} \quad (14)$$

where θ_t is the most recent entry in θ^t.[17]

A full proof is given in the appendix. In the separable case we clearly collapse down to the usual inverse Euler condition, since $\alpha = 0$. The innovation here is to provide a general expression for the marginal cost of incentive-compatible utility provision that also applies under non-separability – at least for the relaxed problem.

The intuition behind (14) is as follows. Changing consumption and output jointly at t for the agent with report history θ^t along the vector $(1, \alpha\left(\theta_t\right))$ increases the within-period utility of that agent at the margin by $u_c\left(\theta_t\right) + u_y\left(\theta_t\right)\alpha\left(\theta_t\right)$ units. By construction it would have the same impact on a mimicking agent with a common report history to $t-1$, but a type one higher at t. The t-dated resource cost of providing utility in this manner at the margin is $1-\alpha\left(\theta_t\right)$ (any extra output being a benefit). Hence the term on the left-hand side of (14) is the marginal cost for every β units of t-dated utility provided, which is converted into $t+1$ resources at the prevailing real interest rate. The term on the right-hand side is, by similar reasoning, the marginal cost of providing the agent with report history θ^t with a guaranteed utility increment of one unit across types at time $t+1$. This uniform utility provision preserves incentive compatibility at t and $t+1$ under the relaxed problem, for reasons that are essentially familiar from the separable case.[18]

Note that, like the 'no distortion at the top' condition, this result applies for general type processes – so long as the first-order approach remains valid.

The marginal cost term that features in (14) will feature frequently in the analysis that follows. For simplicity we will often refer to it as $MC\left(c, y; \theta\right)$:

$$MC\left(c, y; \theta\right) := \frac{1-\alpha\left(c, y; \theta\right)}{u_c\left(c, y; \theta\right)+u_y\left(c, y; \theta\right)\alpha\left(c, y; \theta\right)} \quad (15)$$

[17] We suppress dependence upon past type reports and current allocations.

[18] See, for instance, Golosov, Tsyvinski and Werning (2006).

5.3 Implications for optimal savings wedges

Our generalised inverse Euler condition can be used to make qualitative statements about optimal tax distortions, and in particular optimal savings wedges. It is well established that under preference separability savings are always deterred at the margin under the optimal allocation, in the sense that the marginal utility of consumption in period t will be too low for the standard consumption Euler equation to hold between t and $t+1$. Agents collectively defer consumption for the purpose of self-insurance, but this makes it harder to incentivise them to produce in later years, since their past savings reduce the desirability of additional earnings.

The presence of the terms in u_y in (14) suggests the *consumption* Euler condition may not be the best focal point in assessing savings wedges in the more general case. But the consumption Euler condition is not the only way to state dynamic optimality under autarky: any combined change in consumption and output at t, coupled with any distribution of the saved (or borrowed) proceeds at $t+1$ between consumption and output is possible, and must not increase utility at an optimum under autarky. In particular, in a world with no taxation the following would hold:

$$\frac{u_c(\theta_t) + u_y(\theta_t)\,\alpha(\theta_t)}{1-\alpha(\theta_t)} = \beta R_{t+1} \sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta^t\right) \frac{u_c(\theta_{t+1}) + u_y(\theta_{t+1})\,\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})} \tag{16}$$

An agent's savings decisions are implicitly being distorted whenever equation (16) does not hold, with saving being discouraged whenever the left-hand side is less than the right. Part of this distortion may operate via the allocation of work effort between the different time periods, but this is no less an intertemporal distortion than a wedge in the consumption Euler condition – particularly as we are not specifying distinct 'income' and 'savings' tax instruments at present.

The useful feature of equation (16) is that we can be more precise about deviations from *this* expression at the optimum than we can about deviations from a consumption Euler equation. Specifically, we have the following.

Proposition 4 ***Deterred savings:*** *For all time periods $t \geq 1$ and for all reporting histories θ^t, if consumption and labour supply are either Edgeworth substitutes or additively separable in preferences then savings will be deterred at the optimum, in the sense that the constrained-optimal allocations (c_t, y_t) and X_{t+1} for the given θ^t*

will satisfy the following condition:

$$\frac{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}{1-\alpha(\theta_t)} \leq \beta R_{t+1} \sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta^t\right) \frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})} \tag{17}$$

with the inequality holding strictly so long as the object $\frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})}$ *varies for different draws of* $\theta_{t+1}\in\Theta$*.*

Proof. If consumption and labour supply are substitutes then $\alpha(\theta_t)<0$, so for the preferences we are focusing on we must always have:

$$\frac{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}{1-\alpha(\theta_t)} > 0 \tag{18}$$

(recalling that $u_y<0$). Thus by Jensen's inequality we have the following:

$$\begin{aligned} &\sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta^t\right) \left[\frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})}\right]^{-1} \\ &\geq \left[\sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta^t\right) \frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})}\right]^{-1} \end{aligned} \tag{19}$$

with a strict inequality provided $\frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})}$ varies for different draws of θ_{t+1}. The left-hand side of (19) is also the right-hand side of equation (14); the inequality in the Proposition then follows from using (14) in (19). ■

Note that we are *not* able to state the result for the case of Edgeworth complements: in that case we cannot rule out the possibility that $\frac{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})}{1-\alpha(\theta_{t+1})}<0$ may hold at the optimum for some values of θ_{t+1}, preventing us from applying Jensen's inequality. In fact, Lemma 4 below implies savings will be deterred under complements if types additionally follow an iid process; but this extension does not seem of much practical importance.

In economic terms, the result suggests the problem of over-saving in the absence of perfect insurance markets carries over directly to the case of substitutes. But we cannot be confident that savings should be deterred if the marginal cost of incentive-compatible utility provision defined in (15) could turn negative. Though that possibility may at first appear unlikely, we cannot rule it out under Markov shock processes and complementarity.[19] This would imply that the deferral of *utility* would not be

[19] Numerical simulations confirm that negative values for MC can indeed arise at an optimum. Details are available on request.

accompanied by the deferral of *resources*, and it is perhaps not surprising that in so distorted an environment the standard resistance to private-sector saving – whether interpreted as resource or utility deferral – need not apply. Since type persistance and complementarity appears empirically plausible, the practical implications of this class of model for optimal savings wedges are perhaps subtler than has been appreciated.

Moving away from its implications for marginal tax wedges, it will also be interesting to consider what our generalised inverse Euler condition implies for the 'immiseration'-type results that emerge in the special case that $R_t \equiv \beta^{-1}$. It is well known (see, for instance, Farhi and Werning (2007)) that under separable preferences almost all agents will see their consumption converge to zero along an optimal path in this case.[20] Again, this result will turn out to generalise to the case of substitutes but not of complements. But unfortunately the proofs rely on other arguments that are still to be established, so we defer a full treatment until Section 8.

6 Intratemporal optimality: the case of iid types

This section develops novel insight into the equity-efficiency trade-off at the heart of the dynamic Mirrleesian model in the simple case that type draws are iid. Specifically, we derive a further set of $N-1$ 'intratemporal' optimality conditions that link allocations across agents with common past type draws. Whilst the iid assumption is clearly a restrictive one, these conditions must apply irrespective of the dynamic setting of the model; they therefore also provide fresh insight into the static problem. They show a simple link between the degree of productive inefficiency the policymaker should be willing to accept for any given agent, and a measure of the relative welfare of those whose types are higher.

6.1 Analytical treatment

For the formal arguments it is useful to define the function $\tau : \mathbb{R}_+^2 \times \Theta \to \mathbb{R}$ by $\tau(c, y, \theta) := 1 + \frac{u_y(c,y,\theta)}{u_c(c,y,\theta)}$. Thus τ is the implicit within-period marginal income tax rate faced by an agent of type θ receiving an allocation (c, y).

A first step towards obtaining the general optimality conditions we are after is the following Lemma, the proof of which is in the appendix:

[20] This 'immiseration' was first demonstrated as a potential property of optimal allocations under asymmetric information in a moral hazard context by Thomas and Worrall (1990).

Lemma 1 *Suppose that type draws are iid across agents and time. Fix a vector* $\nu \in \mathbb{R}^N$ *(whose nth element is denoted* ν_n*) such that:*

$$\sum_{n=1}^{N} \pi_\Theta\left(\theta^n\right) \nu_n = 0$$

For all time periods $t \geq 1$ *and any prior reporting history* $\theta^{t-1} \in \Theta^{t-1}$ *it is possible to perturb the set of optimal allocations* $X_t\left(\theta^t\right)$ *in a manner that will preserve the incentive-compatibility constraints of the relaxed problem in all periods whilst raising the within-period utility of an agent of type* θ_t^n *by an amount* $\nu_n \delta$ *in period* t*, for any scalar* δ *satisfying* $|\delta| < \varepsilon$ *for some* $\varepsilon > 0$*.*

Thus any incremental vector that delivers zero *expected* utility across types from the perspective of previous periods can be engineered through changes to within-period allocations alone, without affecting incentive compatibility at any horizon. Our aim is to characterise the impact that perturbations of this kind have on resources at the margin as the scalar δ is moved away from zero. Again, if we start at an optimum this marginal cost must itself be zero.

It turns out that an important object when assessing the marginal resource cost of any perturbation is the specific marginal cost to the policymaker of a movement along the nth agent's within-period indifference curve by an amount just sufficient to reduce by a unit the utility that could be obtained by the $n+1$th agent when mimicking the nth, normalised by the distance between these two types. This can be interpreted as the marginal cost to the policymaker of inducing additional productive distortions into the economy, in order to reduce the mimicking rents that are obtainable by higher-type agents. We label it $DC\left(c, y, \theta\right)$ (the 'distortion cost'):

$$DC\left(c, y; \theta\right) := \frac{\tau\left(c, y; \theta\right)\left[\theta' - \theta\right]}{u_c\left(c, y; \theta'\right)\left(1 - \tau\left(c, y; \theta\right)\right) + u_y\left(c, y; \theta'\right)} \tag{20}$$

where $\theta' = \min\left\{\theta'' \in \Theta : \theta'' > \theta\right\}$.[21] Ultimately the entire set of necessary optimality conditions that we derive will consist of relationships between the objects $MC\left(c, y; \theta\right)$ and $DC\left(c, y; \theta\right)$, expressed at different horizons and for different type draws.

To provide some interpretation, the object in the numerator of (20) is the effective marginal tax rate levied on the agent of type θ – which is what the policymaker

[21] If θ is the maximal element in Θ we can arbitrarily define $u_c\left(c, y; \theta'\right) = 1$ and $u_y\left(c, y; \theta'\right) = 0$ irrespective of c and y. This is for completeness only: we already know that there will be no distortion at the top, in the sense that $\tau\left(c, y; \theta\right) = 0$, in this case. It therefore makes sense to fix $DC\left(c, y; \theta\right)$ to zero too.

foregoes at the margin for every unit by which the production of that agent is reduced – multiplied by the distance between the two relevant types. The denominator is the number of units by which the utility of mimicking types is changed for every unit increase in production for those reporting θ, given that the perturbation takes place along the indifference curve of type θ, whose slope is equal to $(1-\tau)$. Its inverse is thus the number of units by which production must be reduced in order to reduce minicking utility by one unit. So overall the expression gives the amount of lost tax revenue for every reduction in a mimicker's utility by $\theta' - \theta$ units.[22]

The best reading of $DC(c, y; \theta)$ is as an efficiency cost of distorting outcomes. The larger it is, the greater are the deviations from full productive efficiency that the policymaker is willing to tolerate for agents of type θ in order to hold down the mimicking rents of higher types. At times we will refer to $DC(c, y; \theta)$ and 'the labour wedge' interchangeably, though strictly $DC(c, y; \theta)$ is this wedge normalised by the impact productive distortions are having on higher types' utility at the margin.

A full consideration of the marginal resource costs associated with general within-period changes in utilities yields the following, the proof of which is in the appendix:

Proposition 5 ***Intratemporal optimality (iid case):*** *Suppose type draws are iid across agents and time. Then for all time periods $t \geq 1$ and all reporting histories θ^t, the constrained-optimal allocation matrix $X_{t+1}(\theta^t)$ satisfies the following condition:*

$$\sum_{\theta^n \in \Theta \setminus \theta^N} \pi_\Theta(\theta_t^n)(\nu_{n+1} - \nu_n)\frac{DC(\theta_t^n)}{\theta_t^{n+1} - \theta_t^n} = \sum_{\theta^n \in \Theta} \pi_\Theta(\theta_t^n)\nu_n MC(\theta_t^n) \tag{21}$$

where ν_n is the nth element of any vector ν that satisfies:

$$\sum_{n=1}^{N} \pi_\Theta(\theta^n)\nu_n = 0$$

The most useful insights from this condition come when one makes specific choices for the ν vector. To this end, suppose we pick some $m \in \{1, ..., N-1\}$ and let $\nu_n = -1$ for all $n \leq m$ and $\nu_n = [\pi_\Theta(\theta > \theta^m)]^{-1} - 1$ for all $n > m$. By construction the *ex-ante* expected value of ν_n is zero, and we can state the following corollary:

Corollary 2 *Suppose type draws are iid across agents and time. Then for all time periods $t \geq 1$ and all reporting histories θ^t, the optimal allocation matrix $X_{t+1}(\theta^t)$*

[22] The normalisation is useful here because the cost of reducing a higher type's utility by a given amount will generally be higher the closer are the two types.

satisfies the following condition for all $m \in \{1, ..., N-1\}$*:*

$$\frac{\pi_\Theta\left(\theta_t^m\right)}{\theta_t^{m+1}-\theta_t^m}\frac{1}{\pi_\Theta\left(\theta_t>\theta_t^m\right)}DC\left(\theta_t^m\right)=E\left[MC\left(\theta_t\right)|\theta_t>\theta_t^m\right]-E\left[MC\left(\theta_t\right)\right] \tag{22}$$

E here is the standard expectations operator, under the unique within-period (iid) distribution for agents with a common type history. (22) gives a relationship between the expected marginal cost of utility provision across *all* types at t with the given history, and the expected marginal cost conditional upon type being higher than θ_t^m. The condition states that the higher is the gap between these two expected marginal costs, the more willing the policymaker should be to distort the productive activity of an agent of type θ_t^m. The cost of providing utility will generally be higher for those whose types are relatively high, as a by-product of the need to provide incentives. (22) therefore says that the more privileged are higher types, the greater is the productive distortion that the policymaker should be willing to tolerate on any given agent. In this sense it provides a classic statement of the 'efficiency-equity trade-off', with variations in $MC\left(\theta_t\right)$ across the type distribution providing a measure of inequality, and $DC\left(\theta_t^m\right)$ a measure of productive inefficiency.[23]

The first two fractions on the left-hand side of (22) are the equivalent of a 'hazard rate' when types are discrete. The higher is the measure of individuals of type θ_t^m relative to those further up in the distribution, the greater will be the costs of distorting the production of this group in order to reduce higher types' rents. The importance of this hazard rate in characterising optimal allocations in static Mirrleesian models has been emphasised since the work of Diamond (1998). It becomes particularly important if one does not place an upper bound on the support of the type distribution, in which case the limiting hazard rate at the top of the distribution is of great importance to the top rate of tax. Here, if we allow $N=\infty$ a non-zero labour wedge at the top follows whenever $DC\left(\theta_t^m\right)>0$ for limiting values of m. According to (22) this should obtain even in the dynamic case provided the hazard rate is bounded above zero.

It is insightful to contrast our approach in analysing the within-period labour wedge with that taken by Golosov, Troshkin and Tsyvinski (2011). These authors show how to adapt the original Hamiltonian approach to the static income tax problem that Mirrlees (1971) took, and thereby arrive at an augmented expression for the

[23] Recall that if preferences were separable with log consumption utility we would have $MC\left(\theta_t\right)=c\left(\theta_t\right)$, so the association between variations in marginal cost and consumption inequality is particularly strong in that case.

within-period labour wedge that is expressed in terms of compensated and uncompensated labour supply elasticities – deliberately to aid comparison with the equivalent expression in Saez (2001). Because their Hamiltonian problem is set up in a manner that requires simultaneous choice over current allocations and future (promised) values this expression for the labour wedge depends in a fairly complex manner on promised values, even in the iid case. This allows the authors to emphasise the additional optimality considerations that are necessary in a dynamic setting relative to the static problem.

Our approach instead highlights the *similarities* between static and dynamic problems. The set of conditions satisfying (22) is just as necessary when dynamics are absent as when they are present: it is the constraints of the problem rather than its optimality conditions *per se* that become more complex when promised utility can be used to incentivise production. In Section 7 we will see that the same basic intuition holds when more general Markov shocks are introduced, but with some additional complication.

6.2 A simple efficiency-equity trade-off

There is a further simple and insightful summary expression can be obtained from Proposition 5 alone. Suppose that we wish to implement some utility perturbation vector ν whose nth element takes the form:

$$\nu_n = \theta^n - \sum_{\theta^m \in \Theta} \pi_\Theta \left(\theta^m\right) \theta^m \tag{23}$$

This clearly satisfies $\sum_{\theta^n \in \Theta} \pi_\Theta \left(\theta^n\right) \nu_n = 0$, so it provides an admissible vector by which we can augment utilities within a time period at the margin, under the maintained iid assumption.

For the specified ν vector, condition (21) then becomes:[24]

$$E\left[DC\left(\theta_t\right)\right] = Cov\left(\theta_t, MC\left(\theta_t\right)\right) \tag{24}$$

where the expectation and covariance operators are taken under the unique distribution π_Θ, across types with a common shock history prior to t.

This is a remarkably clear statement of the efficiency-equity trade-off that the Mirrleesian model requires. The term on the left-hand side gives the average value of the marginal resource cost the policymaker is willing to endure in order to hold down

[24] Recall that 'no distortion at the top' implies $DC\left(\theta_t^N\right) = 0$.

the utility rents that are enjoyed by high-type agents: higher values for $DC(\theta_t^n)$ imply higher productive distortions. The term on the right-hand side tracks the degree of cross-sectional inequality in the economy: the covariance between θ_t^n and $MC(\theta_t^n)$ will be greater the greater is the degree of inequality in welfare across θ_t^n draws. So (24) states that inequality should be tolerated in proportion to the resource costs of reducing it.

6.3 Optimal effective income tax rates

The results of the earlier analysis also allow us to demonstrate a further quite general result with important economic implications, which does not in fact require the iid assumption on type draws. The proof is in the appendix.

Proposition 6 ***Non-negative income taxes:*** *For all time periods $t \geq 1$, all reporting histories θ^{t-1} and all $\theta_t^n \in \Theta$ the implicit marginal tax rate $\tau(\theta_t^n)$ satisfies $\tau(\theta_t^n) \geq 0$.*

So unlike the savings distortion the direction of the intratemporal distortion on production is completely unambiguous: the optimal effective marginal income tax rate is never negative. In a sense the result should not be surprising. Since a 'downwards' movement along the within-period indifference curve of lower types reduces the utility of higher-type mimickers, it is always better to move to a point where this indfference curve has a slope $(\frac{dc}{dy})$ that is less than one – accepting some productive distortion as the cost of greater equity. Subsidising work would require still higher utility to be granted further up the type distribution, whilst incurring a positive cost.

7 Optimal distortions with persistent types

Whilst the iid model is instructive it is plainly unrealistic as a description of the way individuals' earnings capacities evolve in practice. Practical implementability requires that we analyse persistence in types.[25] The simplest way to do this is to assume the productivity measure π_Θ incorporates a general Markov structure, so that $\pi_\Theta(\theta_{t+1}|\theta^t) = \pi_\Theta(\theta_{t+1}|\theta_t)$.

This gives us an extra complication. For agents with reporting history θ^t we may be able to define a perturbation to allocations at $t+1$ that has zero impact on the

[25] A far more substantive generalisation would be to allow for active human capital accumulation. Kapička (2011b) provides important work in this vein.

expected utility at t of a relevant *truth-telling* agent, but the probability distribution under which this expectation is calculated is now particular to that agent. An agent who is, at the optimum, on the cusp of *falsely* reporting θ^t will take expectations of the future returns from a mimicking strategy under a different probability distribution to the truth-teller – and thus may well experience a change in expected utility even though the truth-teller does not. This would undermine incentive compatibility at t. For this reason the local incentive compatibility constraints become relatively hard to satisfy in the Markov case.

But note that some of the earlier results will go through unchanged. In particular, the Markov and iid cases will be equivalent to one another when we consider perturbations to the allocations at $t+1$ and (possibly) t of any agent whose allocation was not 'envied' at t. This could either be because $t+1$ is the first period of the model, or because the agent's type was the highest possible at t. Proposition 5 will then extend directly to these cases, and the optimality conditions set out in Corollary 2 will again apply. What remains is to understand how optimality conditions are affected when agents' prior allocations *were* envied.

We show that there are two important ways in which optimality requirements change in this case. First, an additional *inter*temporal condition arises, ensuring that the cost to the policymaker of preventing mimicking is spread optimally through time. This condition determines the optimal dynamics for productive distortions, and generalises the recent results of Farhi and Werning (2011) relating to the dynamics of the labour wedge. Second, and offsetting this, we lose one of the intertemporal conditions: equation (22) holds in a 'first-differenced' form only. We consider these points in turn.

7.1 Intertemporal optimality: the dynamics of productive distortions

The richer dynamics introduced into the model when types are persistent have already been given much attention in the work of Farhi and Werning (2011). Here we generalise their study of the dynamics of the labour wedge, and explain why it is usually desirable to give greater emphasis to equity concerns relative to efficiency as time progresses.

Before stating the main argument we must provide an equivalent to Lemma 3 to confirm incentive compatibility for the types of dynamic perturbations we will consider. We have the following, the proof of which is in the appendix:

Lemma 2 *For all time periods $t \geq 1$, all reporting histories θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, and any vector ν that satisfies the two restrictions:*

$$\sum_{\theta^m \in \Theta} \pi_\Theta \left(\theta_{t+1}^m | \theta_t^n\right) \nu_m = 0$$

and

$$\sum_{\theta^m \in \Theta} \pi_\Theta \left(\theta_{t+1}^m | \theta_t^{n+1}\right) \nu_m = 1$$

it is possible to perturb the constrained-optimal allocations $(c_t(\theta^t), y_t(\theta^t))$ and $X_{t+1}(\theta^t)$ in a manner that will preserve the incentive compatibility constraints of the relaxed problem in all periods whilst raising the within-period utility of an agent of type θ_{t+1}^m by an amount $\nu_m \delta$ at $t+1$ for any δ satisfying $|\delta| < \varepsilon$ for some $\varepsilon > 0$ and leaving equilibrium utility in all other periods constant.

This result immediately takes us to the additional dynamic condition that we desire. The proof is in the appendix.

Proposition 7 ***Dynamic cost-spreading:*** *For all time periods $t \geq 1$ and any reporting history θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, the constrained-optimal $t+1$ allocation matrix $X_{t+1}(\theta^t)$ together with the constrained-optimal period-t allocation pair $(c_t(\theta^t), y_t(\theta^t))$ must satisfy the following condition:*

$$\beta R_{t+1} \frac{DC\left(\theta_t^n\right)}{\theta_t^{n+1} - \theta_t^n} = \sum_{\theta^m \in \Theta \backslash \theta^N} \pi_\Theta \left(\theta_{t+1}^m | \theta_t^n\right) \left(\nu_{m+1} - \nu_m\right) \frac{DC\left(\theta_{t+1}^m\right)}{\theta_{t+1}^{m+1} - \theta_{t+1}^m} - \sum_{\theta^m \in \Theta} \pi_\Theta \left(\theta_{t+1}^m | \theta_t^n\right) \nu_m MC\left(\theta_{t+1}^m\right) \tag{25}$$

where ν_m is the mth element of any vector ν that satisfies the two restrictions given in Lemma 2.

In general this condition seems likely to result in greater *conditional* equality at $t+1$ the higher is the productive distortion for an agent at t (measured by $DC(\theta)$) – where equality here is considered across agents with a common report history up to t. If it is worthwhile to induce productive distortions at t to reduce mimicking rents then it should also be worthwhile reducing utility disparities across types at $t+1$ to the same end: 'twisting' utility allocations in this way is to the detriment of would-be mimickers in period t, whose assessment about their period-$t+1$ type draw is relatively optimistic by comparison with truth-tellers.[26]

The next sub-section clarifies further the sense in which productive efficiency is given less and less weight through time, and in the process provides a novel expression for the dynamics of the labour wedge, and of the policymaker's willingness to trade off equity and efficiency through time.

7.2 Equity, efficiency, and the dynamics of the labour wedge

To extract a more interpretable expression from (25) we can fix a utility increment vector ν for application at $t+1$ whose mth entry is given by:

$$\nu_m = \frac{\left\{\theta_{t+1}^m - \sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta_t^n\right)\theta_{t+1}\right\}}{\sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta_t^{n+1}\right)\theta_{t+1} - \sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta_t^n\right)\theta_{t+1}} \tag{26}$$

The denominator here is the difference in expected type in period t between those who drew θ_t^{n+1} and those who drew θ_t^n in the previous period (with $n < N$), whilst the numerator is the difference between θ_{t+1}^m and its expected value for the previous period's truth-tellers. It is clear by inspection that this vector will satisfy the requirements of a ν vector in Proposition 7. Re-writing condition (25) then gives:

$$\beta R_{t+1} DC\left(\theta_t^n\right) = \frac{E\left[DC\left(\theta_{t+1}\right)|\theta_t^n\right] - Cov\left[\left(\theta_{t+1}, MC\left(\theta_{t+1}\right)\right)|\theta_t^n\right]}{\left\{E\left[\theta_{t+1}|\theta_t^{n+1}\right] - E\left[\theta_{t+1}|\theta_t^n\right]\right\}/\left(\theta_t^{n+1} - \theta_t^n\right)} \tag{27}$$

where the conditional expectation and covariance terms are again taken across types with common productivity draws to t.[27]

The term on the left-hand side of the equation is a measure of the marginal efficiency cost of the labour tax wedge imposed in period t on type θ_t^n. Proposition 6 implies this object will always be non-negative. The term on the right-hand side can be digested in pieces. The fraction's numerator is familiar from equation (24) above: it gives the difference between the average of the marginal efficiency costs of within-period tax distortions implemented in period $t+1$, and the costs from excessive

[26] Golosov, Troshkin and Tsyvinski (2011) provide similar intuition in accounting for an augmented weighting structure when controlling promised utilities in the Hamiltonian problem that they study. This places more weight on utility provided in states of the world for which there is a relatively large difference in the probability of occurrence between truth-tellers and mimickers from the previous period. The authors identify this as a more redistributionary force, though their analytical approach is very different from ours: they do not provide a direct summary relationship comparable to equation (27) below.

[27] Farhi and Werning (2011) provide a special case of this condition under the assumptions that consumption and labour supply are separable in preferences, the disutility of work is isoelastic, and productivity is AR(1).

inequality. The denominator, meanwhile, is the difference in the expected type draw for $t+1$ between agents whose t-period draws are θ_t^{n+1} and θ_t^n respectively (relative to the absolute difference between these types). It is a natural measure of the persistence of productivity draws.

Taken together, (27) has a clear message: persistent shocks mean the policymaker should implement more equality for any given within-period efficiency cost. This is what it means for the numerator of the fraction to be positive, which – according to (27) – it must be under any optimal scheme whenever the tax wedge was strictly positive in the prior period. Note that this relatively high equality only holds across agents with a common shock history up to period t; it need not carry over to the economy as a whole. Nonetheless, it seems likely to be behind the simulation result obtained under separable preferences by Farhi and Werning (2011) that optimal tax wedges drift upwards on average over time when log type draws are I(1).

Interestingly Farhi and Werning focus on the cross-sectional profile of the labour wedge to argue that the extra dynamic complications introduced by (27) result in more *regressive* outcomes relative to the iid case – since marginal taxes tend to drift upwards most over time for those who receive relatively low type draws. The interpretation we provide here appears quite the opposite: persistent shocks introduce a bias towards equity relative to efficiency. But this is because high marginal tax rates at any given point in the income distribution are a way for average rates to be increased higher up. The reason for inducing within-period labour supply distortions in the first place is to reduce the compensation that more productive agents must be paid to incentivise them to work. Hence the regressivity in marginal rates identified by Farhi and Werning is really just a means for engineering greater *progressivity* in utility outcomes in equilibrium: the two interpretations are not inconsistent.

Two other interesting observations can be made about (27). First, note that it has a 'reset' feature. If at any point in time t the highest productivity θ_t^N is drawn then the set of optimality conditions at $t+1$ becomes identical to the iid case. This implies in particular that equation (24) will hold: the numerator of the fraction on the right-hand side of (27) will equal zero, with efficiency and equity considerations exactly balancing. So any average drift towards greater equality always has the potential to be dominated by 'no distortion at the top'.

Second, if type draws are I(1) and the interest rate is equal to the inverse of the discount factor then the object $DC(\theta)$ will in general increase over time. In this event the denominator of the fraction on the right-hand side is one, and since the covariance term will be (at least weakly) positive the equation reduces to a random walk plus a stochastic drift term. Thus full persistence in type draws translates into persistence *plus drift* in labour wedges. This is a generalisation of the 'tax smoothing' result of Golosov, Tsyvinski and Werning (2006), who showed that labour wedges should be constant through time in the event that type draws are drawn 'once and for all' in the first period of the model. In that case there is no variation in type draws at $t+1$ for any given draw at t, implying the covariance term drops from (27) whilst the conditional expectations reduce to certainties. Allowing continued uncertainty instead biases us away from tax smoothing and towards continued upwards drift in the labour tax wedge.

7.3 Intratemporal optimality: a dimension lost

If we are considering a perturbation that applies exclusively in period $t+1$ to the allocations of agents with a common reporting history θ^t, such that $\theta_t \neq \theta_t^N$, we need to make sure that this perturbation does not affect the incentive at t for truthful reporting – either for an agent whose true type is θ_t^n or for one whose true type is θ_t^{n+1} (and thus is tempted to mimic θ_t^n). This implies that the expected utility consequences of the perturbation must be zero under both the 'truth-teller's' probability measure $\pi_\Theta\left(\cdot|\theta_t^n\right)$ and the 'mimicker's' measure $\pi_\Theta\left(\cdot|\theta_t^{n+1}\right)$. In the iid case we were able implement any vector of marginal utility increments across agents at $t+1$ provided this vector satisfied $\sum_{\theta^n\in\Theta}\pi_\Theta\left(\theta^n\right)\nu_n = 0$ for the *common* measure across $t+1$ types, π_Θ.

When shocks are Markov the probability measure across $t+1$ types is no longer common. But we can preserve incentive compatibility for both truth-tellers and relevant mimickers provided we perturb utilities at the margin according to a vector ν that *jointly* satisfies the *two* conditions:

$$\sum_{\theta^m\in\Theta}\pi_\Theta\left(\theta_{t+1}^m|\theta_t^n\right)\nu_m = \sum_{\theta^m\in\Theta}\pi_\Theta\left(\theta_{t+1}^m|\theta_t^{n+1}\right)\nu_m = 0 \tag{28}$$

In general one can always find $N-2$ linearly independent ν vectors for which this condition is satisfied, as against $N-1$ that satisfy the single restriction in the iid case. Thus the movement to Markov probabilities denies us the capacity to carry out intratemporal perturbations in precisely one dimension.

Lemma 1 can now be easily adjusted to cover intratemporal perturbations in the Markov case:

Lemma 3 *For all time periods $t\geq 1$, all reporting histories θ^t such that $\theta_t = \theta_t^n \neq$*

θ_t^N, and any vector ν that satisfies (28) it is possible to perturb the constrained-optimal allocations $X_{t+1}(\theta^t)$ in a manner that will preserve the incentive compatibility constraints of the relaxed problem in all periods whilst raising the within-period utility of an agent of type θ_{t+1}^n by an amount $\nu_n\delta$ at $t+1$ for any δ satisfying $|\delta| < \varepsilon$ for some $\varepsilon > 0$ and leaving utility in all other periods constant.

We omit a proof, since the logic is identical to that of Lemma 1, except that it is applied here only to the subset of within-period perturbations admissible in the Markov case. The required intratemporal optimality conditions can then be stated formally:

Proposition 8 ***Intratemporal optimality (Markov case):*** *For all time periods $t \geq 1$ and any reporting history θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, the constrained-optimal allocation $X_{t+1}(\theta^t)$ satisfies the following condition:*

$$\sum_{\theta^m \in \Theta \backslash \theta^N} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^n\right)\left(\nu_{m+1} - \nu_m\right) \frac{DC\left(\theta_{t+1}^m\right)}{\theta_{t+1}^{m+1} - \theta_{t+1}^m} = \sum_{\theta^m \in \Theta} \pi_\Theta\left(\theta_{t+1}^m | \theta_t^n\right) \nu_m MC\left(\theta_{t+1}^m\right) \tag{29}$$

where ν_m is the mth element of any vector ν that satisfies (28).

The proof again repeats earlier arguments so is omitted. The additional restriction on ν in (28) relative to the single requirement imposed for Proposition 5 essentially limits us to a 'first-differenced' equivalent to Corollary 2, which is insightful in spite of its unwieldiness:

Corollary 3 *For all time periods $t \geq 1$ and any reporting history θ^t such that $\theta_t = \theta_t^n \neq \theta_t^N$, the constrained-optimal allocation $X_{t+1}(\theta^t)$ satisfies the following condition for all $m \in \{1, ..., N-2\}$:*

$$\begin{aligned} &\frac{\left\{\frac{\pi_\Theta\left(\theta_{t+1}^m|\theta_t^n\right)}{\theta_{t+1}^{m+1}-\theta_{t+1}^m}\frac{1}{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^m|\theta_t^n\right)}DC\left(\theta_{t+1}^m\right) - \left[E\left[MC\left(\theta_{t+1}\right)|\theta_{t+1}>\theta_{t+1}^m\right] - E\left[MC\left(\theta_{t+1}\right)\right]\right]\right\}}{\left[\frac{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^m|\theta_t^{n+1}\right)}{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^m|\theta_t^n\right)} - 1\right]} \\ &= \frac{\left\{\frac{\pi_\Theta\left(\theta_{t+1}^{m+1}|\theta_t^n\right)}{\theta_{t+1}^{m+2}-\theta_{t+1}^{m+1}}\frac{1}{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^{m+1}|\theta_t^n\right)}DC\left(\theta_{t+1}^{m+1}\right) - \left[E\left[MC\left(\theta_{t+1}\right)|\theta_{t+1}>\theta_{t+1}^{m+1}\right] - E\left[MC\left(\theta_{t+1}\right)\right]\right]\right\}}{\left[\frac{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^{m+1}|\theta_t^{n+1}\right)}{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^{m+1}|\theta_t^n\right)} - 1\right]} \end{aligned} \tag{30}$$

where the expectations are taken using the distribution of agents truthfully reporting θ_t^n in period t, across types with the common history θ^t.

In the iid case we had a within-period equity-efficiency trade-off at each point in the type distribution, linking the optimal degree of productive distortion at that point to the comparative welfare of higher types – as captured by the difference between conditional and unconditional expectations of $MC(\theta)$. Here this is replaced by a relationship between *deviations* from that earlier efficiency condition, which would have set the numerator of each of the main fractions in (30) to zero. Moreover, it follows straightforwardly from Propositions 6 and 7 that the numerator of the main fractions in (30) will be positive whenever the type distribution satisfies a first-order stochastic dominance property – so that $\pi_\Theta\left(\cdot|\theta_t^{n+1}\right)$ first-order stochasically dominates $\pi_\Theta\left(\cdot|\theta_t^{n}\right)$.[28] That is, persistence provides a reason to give greater concern to equity relative to efficiency, so that $DC(\theta)$ is elevated relative to the difference between conditional and unconditional expectations over $MC(\theta)$.

The Corollary implies that, all things equal, this greater weight on equity should be most enhanced in regions of Θ for which the conditional distributions $\pi\left(\cdot|\theta_t^{n}\right)$ and $\pi\left(\cdot|\theta_t^{n+1}\right)$ differ the most – that is, for which the denominator in the main fraction is high. This again derives from a concern to minimise the rents from past mimicking: those of type θ_t^{n+1} will have less of an incentive to report θ_t^{n} the lower are the future returns they can expect from this strategy. This makes the policymaker particularly content to restrict welfare at $t+1$ across regions in the type space to which the distribution conditional on θ_t^{n+1} attaches relatively large weight. To do this requires additional productive distortions at values for θ_{t+1}^{m} for which the ratio $\frac{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^{m}|\theta_t^{n+1}\right)}{\pi_\Theta\left(\theta_{t+1}>\theta_{t+1}^{m}|\theta_t^{n}\right)}$ is high, since these productive distortions are the means by which the mimicking rents of types above θ_{t+1}^{m} can be reduced.[29]

8 Martingale convergence results

The final major area on which we focus is the evolution of optimal outcomes over time, and in particular at the limit as the time horizon becomes large. Suppose that the real interest rate were in all time periods equal to the inverse of the discount

[28]To see this note that when first-order stochastic dominance holds a scalar multiple of the ν vector used to obtain Corollary 2 will be admissible under Proposition 7. The result then follows from the fact productive distortions are non-negative – an immediate implication of Proposition 6.

[29]Exploiting differences in distributions between types in this fashion shares interesting parallels with the work of Crémer and Mclean (1988) in the domain of auction theory. These authors showed *all* information rents could potentially be eliminated by an auctioneer who exploited variations in probability assessments across agents with distinct valuations. Our optimal tax problem does not assume quasi-linear utility, which contributes to the strength of the Crémer-Mclean result, but the intuition otherwise appears very similar.

factor β. Then the generalised inverse Euler equation can be written as:

$$\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} = \sum_{\theta_{t+1}\in\Theta} \pi_\Theta(\theta_{t+1}|\theta_t) \frac{1-\alpha(\theta_{t+1})}{u_c(\theta_{t+1})+u_y(\theta_{t+1})\alpha(\theta_{t+1})} \tag{31}$$

That is to say, we have a martingale in $MC(\theta)$, which we have chosen to write out in full here. When preferences are separable between consumption and labour supply, $\alpha(\theta_t) = 0$ holds, and the expression collapses to a martingale in the inverse of the marginal utility of consumption – an object that is bounded below at zero. As many authors have observed, this boundedness allows the application of Doob's martingale convergence theorem, which implies almost-sure convergence in the inverse marginal utility of consumption to a finite (possibly random) limit. If one can also show that the optimum will never involve consumption staying fixed at a non-zero value, convergence to zero consumption – 'immiseration' – becomes the only possibility.

To generalise these results to the case at hand we need to put a bound on the object in (31) for preference structures beyond the separable case. When consumption and labour supply are Edgeworth substitutes this is straightforward. But when they are Edgeworth complements our scope for doing so is limited. Taken together we have the following result.

Lemma 4 $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} > 0$ *always holds under an optimal plan that solves the restricted problem whenever consumption and labour supply are (a) Edgeworth substitutes, or (b) additively separable. Additionally if consumption and labour supply are Edgeworth complements* $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} > 0$ *will hold when type draws are iid.*

Proof. With separability $\alpha(\theta_t) = 0$, and $u_c(\theta_t) > 0$ is enough. When consumption and labour supply are Edgeworth substitutes we have $\alpha(\theta_t^n) < 0$, and the result is again trivial (recalling $u_y < 0$). When consumption and labour supply are Edgeworth complements and type draws are iid the reasoning is far more involved, and we relegate it to an appendix. ■

We note in passing that this Lemma allows a slight strengthening of Proposition 4, which goes through whenever $MC(\theta)$ is bounded below at zero: we can now assert that it will be optimal to deter savings (in the sense used in that Proposition) in the case of complements *if* type draws are iid.

Having put a zero lower bound on the marginal cost of utility provision for these specific cases, when $R_t = \beta^{-1}$ for all t a direct application of Doob's martingale convergence theorem implies the object $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}$ must converge almost surely along all realisations of θ^∞ to some value $X \in [0, \infty)$, where X is potentially a random variable. We want to be able to say more about the value of X. In fact, it turns out – as in the separable case – that X must equal zero. The next Proposition establishes this.

Proposition 9 ***Convergence:*** *Suppose $R_t = \beta^{-1}$ for all $t \geq 1$. Then $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} \overset{a.s.}{\to} 0$ holds under any optimal plan that solves the restricted problem unless consumption and labour supply are Edgeworth complements and productivities follow a non-iid process.*

Proof. See appendix. ∎

This result is an obvious generalisation of the 'immiseration' results obtained by studying convergence of the standard inverse Euler condition. Moreover, almost sure *consumption* immiseration is a direct implication of this result when one recalls that $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)} = \frac{1}{u_c(\theta_t)}$ when $\theta_t = \theta_t^N$ (the highest type): the outcome for an agent who draws the top productivity parameter in the tth period *must* be zero consumption (almost surely) at the limit as t becomes large, and incentive compatibility then demands that all lower types with the same history must have a still worse lot. So the more complicated nature of the expression for the marginal cost of utility provision in the non-separable case does not undermine the extreme predictions regarding long-run consumption when martingale convergence *can* be applied.

Perhaps the more surprising implication of this section, though, is that when productivity follows a Markov process and consumption and labour supply are Edgeworth complements – so that those who are working longer hours with a given level of consumption have a higher marginal utility of consumption – we *cannot* put a zero lower bound on the marginal cost of utility provision. Indeed, this marginal cost may turn negative. We have been able to confirm this through simple numerical simulations of a finite-horizon model,[30] in which the choices of low types are, at the optimum, distorted sufficiently far away from a point of productive efficiency that even movement along a vector giving *equal* consumption and output increments would still raise their utility by more than it would raise the utility of higher-type mimickers, owing to the strength of complementarities. Thus output must be increased by *more* than consumption at the margin to provide balanced utility increments to mimickers and truth-tellers.

[30] Details available on request.

9 Conclusion

This paper shows that substantial original insight into a dynamic optimal income tax problem of the type studied in the New Dynamic Public Finance literature can be obtained by applying a carefully-chosen set of perturbations to optimal allocations, under the assumption that the 'first-order approach' to incentive compatibility is valid. Our main results are as follows.

First, we have been able to provide a novel interpretation of the 'equity-efficiency' trade-off in the class of models considered, showing how the expected value of the resource cost from labour supply distortions should be linked to a measure of the covariance between agents' types and the marginal cost of providing utility to them in an incentive-compatible manner. Since this marginal cost will generally be higher the greater the quantity of utility an agent already enjoys, this covariance term is interpretable as a measure of inequality. In the simple case of iid productivities it is best to set this equity measure equal to the expected value of productive distortions across agents with a common type history.

More surprisingly, we also show that when productivities are *persistent* through time the policymaker should be willing to tolerate more and more productive inefficiency across types with a common history, relative to the degree of inequality. The associated optimality condition suggests that provided persistence is great enough, effective marginal tax rates should drift upwards through time on average. This result has already been obtained numerically, and analytically under specific preference assumtions, by Farhi and Werning (2011). Here we show the extent to which it will generalise, and in particular show that there is an incentive to raise productive distortions most at those points in the type distribution where there is the greatest disparity between distribution functions of agents whose types were adjacent in the *preceeding* time period. These results derive from an incentive to reduce the mimicking rents of higher types by exploiting differences in their distribution functions relative to those that they would mimic.

Turning to optimal savings taxes, it is now well known that in the event of separability between consumption and labour supply it is optimal to apply a positive tax wedge to savings, in the sense that the marginal utility of consumption in period t is below its expected value at $t+1$ (allowing for discounting and the interest rate); this follows from the well-known 'inverse Euler equation'. We have been able to generalise this result, and confirm that savings are always deterred at the optimum (in an economically meaningful sense) unless consumption and labour supply are Edgeworth complements *and* productivity draws are non-iid – but in the case of persistence and complementarity an effective marginal subsidy to savings could not be ruled out.

This latter result has strong connections with a further area that we have investigated: allocations in the long run. Once again, except in the case that consumption and labour supply are Edgeworth complements and productivity draws are non-iid, we have been able to put a zero lower bound on the marginal cost of incentive-compatible utility provision – which in turn will follow a martingale process in the event that the real interest rate equals the inverse of the discount factor β. Martingale convergence theorems then imply almost-sure immiseration for all agents in the economy. But with complementarity and Markov shocks the marginal cost of utility provision may turn negative, and so immiseration need not take place.

Together these results frame as an important empirical question the exact nature of consumers' labour supply-consumption preference structure, as well as the degree to which earnings capacities are persistent. The work of Basu and Kimball (2002) certainly suggests complementarity is the empirically relevant case. If this is indeed true then the results of this paper imply two of the most characteristic predictions of the New Dynamic Public Finance literature under separability – that savings should be deterred and that long-run immiseration characterises optimal allocations – need to be strongly qualified.

10 References

References

[1] Basu, S., and M.S. Kimball (2002), 'Long-run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption', Manuscript, University of Michigan.

[2] Broer, T., M. Kapička and J. Klein (2011), 'Consumption Risk Sharing under Private Information when Earnings are Persistent', Manuscript (available at http://paulklein.se/newsite/research/risksharing.pdf).

[3] Crémer, J., and R.P. McLean (1988), 'Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions', *Econometrica*, 56, 1247–1257

[4] Dasgupta, P. (1982), 'Utilitarianism, Information and Rights', in A. Sen and B. Williams (eds.), *Utilitarianism and Beyond*, Cambridge, Cambridge University Press.

[5] Diamond, P.A., 'Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates', *American Economic Review*, 88, 83–95.

[6] Diamond, P.A., and J.A. Mirrlees (1978), 'A Model of Social Insurance with Variable Retirement', *Journal of Public Economics*, 10, 295–336.

[7] Diamond, P.A. and E. Saez (2011), 'The Case for a Progressive Tax: From Basic Research to Policy Recommendations', *Journal of Economic Perspectives*, 25, 165–190.

[8] Farhi, E. and I. Werning (2007), 'Inequality and Social Discounting', *Journal of Political Economy*, 115, 365–402.

[9] Farhi, E. and I. Werning (2011), 'Insurance and Taxation over the Life Cycle', Manuscript (available at http://econ-www.mit.edu/files/6588).

[10] Golosov, M., N. Kocherlakota and A. Tsyvinski (2003), 'Optimal Indirect and Capital Taxation', *Review of Economic Studies*, 70, 569–587.

[11] Golosov, M., M. Troshkin and A. Tsyvinski (2011b), 'Optimal Dynamic Taxes', Manuscript (available at http://scholar.princeton.edu/golosov/files/odt30.pdf)

[12] Golosov, M., A. Tsyvinski and I. Werning (2006), 'New Dynamic Public Finance: A User's Guide', in *NBER Macroeconomic Annual 2006*, Cambridge MA, MIT Press.

[13] M. Kapička (2011a), 'Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach', Manuscript (available at http://www.econ.ucsb.edu/~mkapicka/persistent.pdf), and forthcoming in *Review of Economic Studies*.

[14] M. Kapička (2011b), 'The Dynamics of Optimal Taxation when Human Capital is Endogenous', Manuscript (available at http://www.econ.ucsb.edu/~mkapicka/OTandHCDynamics.pdf).

[15] King, R.G., C.I. Plosser and S.T. Rebelo (1988), 'Production, Growth and Business Cycles:. The Basic Neoclassical Model', *Journal of Monetary Economics*, 21, 195–232.

[16] Kocherlakota, N.R. (2005), 'Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation', *Econometrica*, 73, 1587–1621.

[17] Kocherlakota, N.R. (2011), *The New Dynamic Public Finance*, Princeton, Princeton University Press.

[18] Mirrlees, J.A. (1971), 'An Exploration in the Theory of Optimum Income Taxation', *Review of Economic Studies*, 38, 175–208.

[19] Pavan, A., I. Segal and J. Toikka (2011), 'Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure', Manuscript (available at http://www.stanford.edu/~isegal/dmd.pdf).

[20] Roberts, K. (2000), 'A Reconsideration of the Optimal Income Tax', in Hammond, P.J., and G.D. Myles (eds.), *Incentives and Organization: Papers in Honour of Sir James Mirrlees* (Oxford: Oxford University Press).

[21] Saez, E. (2001), 'Using Elasticities to Derive Optimal Income Tax Rates', *Review of Economic Studies*, 68, 205–229.

[22] Thomas, J.P. and T. Worrall (1990), 'Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Model', *Journal of Economic Theory*, 51, 367–390.

A Appendix

A.1 Proof of Proposition 1

For the sake of clarity we index the N elements of Θ in ascending order, so $\theta_t^n > \theta_t^m$ whenever $n > m$ for all $n, m \in \{1, ..., N\}$. We have imposed that

$$W\left(\theta_t^n; \theta_t^n, \sigma^{t-1}\right) = W\left(\theta_t^{n-1}; \theta_t^n, \sigma^{t-1}\right)$$

for all $n \in \{2, ..., N\}$, and wish to show that this implies

$$W\left(\theta_t^n; \theta_t^n, \sigma^{t-1}\right) \geq W\left(\theta_t^m; \theta_t^n, \sigma^{t-1}\right)$$

for all $m \in \{1, ..., N\}$, given the increasing differences condition

We first consider the case in which $n > 1$, and show

$$W\left(\theta_t^n; \theta_t^n, \sigma^{t-1}\right) \geq W\left(\theta_t^m; \theta_t^n, \sigma^{t-1}\right)$$

for all $m \in \{1, ..., n-1\}$. For $m = n-1$ this holds by assumption. For $m = n-2$ we have by increasing differences:

$$\begin{aligned} & W\left(\theta_t^{n-1}; \theta_t^n, \sigma^{t-1}\right) - W\left(\theta_t^{n-1}; \theta_t^{n-1}, \sigma^{t-1}\right) \\ & > W\left(\theta_t^{n-2}; \theta_t^n, \sigma^{t-1}\right) - W\left(\theta_t^{n-2}; \theta_t^{n-1}, \sigma^{t-1}\right) \end{aligned}$$

But

$$W\left(\theta_t^{n-1}; \theta_t^n, \sigma^{t-1}\right) = W\left(\theta_t^n; \theta_t^n, \sigma^{t-1}\right)$$

and

$$W\left(\theta_t^{n-2}; \theta_t^{n-1}, \sigma^{t-1}\right) = W\left(\theta_t^{n-1}; \theta_t^{n-1}, \sigma^{t-1}\right)$$

so the prior inequality implies

$$W\left(\theta_t^n; \theta_t^n, \sigma^{t-1}\right) > W\left(\theta_t^{n-2}; \theta_t^n, \sigma^{t-1}\right)$$

as required. Taking $m = n-3$, we then have by increasing differences:

$$\begin{aligned} & W\left(\theta_t^{n-2}; \theta_t^n, \sigma^{t-1}\right) - W\left(\theta_t^{n-2}; \theta_t^{n-2}, \sigma^{t-1}\right) \\ & > W\left(\theta_t^{n-3}; \theta_t^n, \sigma^{t-1}\right) - W\left(\theta_t^{n-3}; \theta_t^{n-2}, \sigma^{t-1}\right) \end{aligned}$$

Again, by

$$W\left(\theta_t^{n-3}; \theta_t^{n-2}, \sigma^{t-1}\right) = W\left(\theta_t^{n-2}; \theta_t^{n-2}, \sigma^{t-1}\right)$$

this inequality collapses to

$$W\left(\theta_t^{n-2}; \theta_t^n, \sigma^{t-1}\right) > W\left(\theta_t^{n-3}; \theta_t^n, \sigma^{t-1}\right)$$

and we can apply the earlier result

$$W\left(\theta_t^n; \theta_t^n, \sigma^{t-1}\right) > W\left(\theta_t^{n-2}; \theta_t^n, \sigma^{t-1}\right)$$

to assert

$$W\left(\theta_t^n; \theta_t^n, \sigma^{t-1}\right) > W\left(\theta_t^{n-3}; \theta_t^n, \sigma^{t-1}\right)$$

as required. The same argument can be applied for all $m \in \{1, ..., n-1\}$.

When $n < N$ we must in the same way consider the cases of $m \in \{n+1, ..., N\}$. For $m = n+1$, we have immediately by the binding restriction on $n+1$-types, together with increasing differences:

$$\begin{aligned} 0 &= W\left(\theta_t^{n+1}; \theta_t^{n+1}, \sigma^{t-1}\right) - W\left(\theta_t^{n}; \theta_t^{n+1}, \sigma^{t-1}\right) \\ &> W\left(\theta_t^{n+1}; \theta_t^{n}, \sigma^{t-1}\right) - W\left(\theta_t^{n}; \theta_t^{n}, \sigma^{t-1}\right) \end{aligned}$$

as required. By similar logic, for $m = n+2$ we have:

$$\begin{aligned} 0 &= W\left(\theta_t^{n+2}; \theta_t^{n+2}, \sigma^{t-1}\right) - W\left(\theta_t^{n+1}; \theta_t^{n+2}, \sigma^{t-1}\right) \\ &> W\left(\theta_t^{n+2}; \theta_t^{n}, \sigma^{t-1}\right) - W\left(\theta_t^{n+1}; \theta_t^{n}, \sigma^{t-1}\right) \end{aligned}$$

and the condition

$$W\left(\theta_t^{n}; \theta_t^{n}, \sigma^{t-1}\right) > W\left(\theta_t^{n+1}; \theta_t^{n}, \sigma^{t-1}\right)$$

then delivers the required result. Again, we can apply an identical argument inductively for all remaining $m < N$. This completes the proof.

A.2 Proof of Proposition 3

Indexing the elements of Θ in ascending order $\{1, ..., N\}$, our strategy is to construct perturbations in both time periods that change the consumption and output levels of the agent reporting θ^n in just such a way that the impact on within-period utility will be identical whether that agent is of true type θ^n or θ^{n+1}. To this end, let $\Delta_{-1}(\delta)$ be given by:

$$\Delta_{-1}(\delta) = \left(\phi^c\left(-\beta\delta; c_t, y_t, \theta_t\right), \phi^y\left(-\beta\delta; c_t, y_t, \theta_t\right)\right) \tag{32}$$

where $\phi^c(k; c, y, \theta)$ and $\phi^y(k; c, y, \theta)$ are defined implicitly when $\theta \neq \max\{\theta'' \in \Theta\}$ by the pair of equalities:

$$u\left(c + \phi^c(k; c, y, \theta), y + \phi^y(k; c, y, \theta); \theta\right) = u(c, y; \theta) + k \tag{33}$$

$$u\left(c + \phi^c(k; c, y, \theta), y + \phi^y(k; c, y, \theta); \theta'\right) = u(c, y; \theta') + k \tag{34}$$

for $\theta' = \min\{\theta'' \in \Theta : \theta'' > \theta\}$, and when $\theta = \max\{\theta'' \in \Theta\}$ by

$$u\left(c + \phi^c(k; c, y, \theta), y; \theta\right) = u(c, y; \theta) + k \tag{35}$$

$$\phi^y(k; c, y, \theta) = 0 \tag{36}$$

That is to say, $\phi^c(k; c, y, \theta)$ and $\phi^y(k; c, y, \theta)$ are the consumption and output increments required to increase the utility of both mimickers and truth-tellers by k units. These functions will be uniquely defined, by the single crossing property. Similarly, the nth row of $\Delta(\delta)$is given by:

$$\left[\phi^c\left(\delta; c_{t+1}, y_{t+1}, \theta_{t+1}^n\right), \phi^y\left(\delta; c_{t+1}, y_{t+1}, \theta_{t+1}^n\right)\right] \tag{37}$$

where we index by type in ascending order. By construction this perturbation must preserve incentive compatibility in the relaxed problem at $t+1$, since the within-period utility that any agent can gain from mimicking a type one lower is being changed by the same amount (δ) as the within-period utility from truth-telling. It must also preserve incentive compatibility at t under the relaxed problem, since its impact on discounted expected utility in period t and earlier is zero, both for agents of type θ_t and for mimickers whose type is one higher. The overall impact of the perturbation on the present value (assessed at time t) of the resources used by the policymaker is given by:

$$\begin{aligned}
&\pi_\Theta\left(\theta^t\right)\left[\phi^c\left(-\beta\delta; c_t, y_t, \theta_t^n\right) - \phi^y\left(-\beta\delta; c_t, y_t, \theta_t^n\right)\right] \\
&+ R_{t+1}^{-1}\pi_\Theta\left(\theta^t\right) \sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta^t\right)\left[\phi^c\left(\delta; c_{t+1}, y_{t+1}, \theta_{t+1}^n\right)\right. \\
&\left.-\phi^y\left(\delta; c_{t+1}, y_{t+1}, \theta_{t+1}^n\right)\right]
\end{aligned}$$

We require for optimality that the derivative of this expression with respect to δ should equal zero when $\delta = 0$; otherwise the policymaker could use fewer resources in obtaining the same aggregate utility. Taking the derivative gives the optimality condition:

$$\begin{aligned}
&\beta\left[\phi_1^c\left(0; c_t, y_t, \theta_t\right) - \phi_1^y\left(0; c_t, y_t, \theta_t\right)\right] \\
&= R_{t+1}^{-1} \sum_{\theta_{t+1}\in\Theta} \pi_\Theta\left(\theta_{t+1}|\theta^t\right)\left[\phi_1^c\left(0; c_{t+1}, y_{t+1}, \theta_{t+1}^n\right)\right. \\
&\left.-\phi_1^y\left(0; c_{t+1}, y_{t+1}, \theta_{t+1}^n\right)\right]
\end{aligned} \tag{38}$$

where ϕ_1^c denotes the derivative of ϕ^c with respect to its first argument. By total differentiation of conditions (33) to (36) with respect to k it is easy to show:

$$\phi_1^c(0; c, y, \theta) - \phi_1^y(0; c, y, \theta) = \frac{1-\alpha(c, y; \theta)}{u_c(c, y; \theta) + u_y(c, y; \theta)\,\alpha(c, y; \theta)} \tag{39}$$

The result follows.

A.3 Proof of Lemma 1

We need to show that it is possible to change the consumption and output levels of each agent in such a way that utilities change in the manner described in the Lemma, and 'downwards' incentive compatibility restrictions remain satisfied for all δ in an open neighbourhood of 0. This requires that the following two conditions are satisfied at t for all $n \in \{1, ..., N\}$:

$$u\left(c_{n,t} + \delta_n^c(\delta), y_{n,t} + \delta_n^y(\delta); \theta_t^n\right) = u\left(c_{n,t}, y_{n,t}; \theta_t^n\right) + \nu_n \delta \tag{40}$$

$$u\left(c_{n,t} + \delta_n^c(\delta), y_{n,t} + \delta_n^y(\delta); \theta_t^{n+1}\right) = u\left(c_{n,t}, y_{n,t}; \theta_t^{n+1}\right) + \nu_{n+1} \delta \tag{41}$$

where $\delta_n^c(\delta)$ and $\delta_n^y(\delta)$ are the perturbations to the nth agent's consumption and output levels respectively. For the Nth agent we just need:

$$u\left(c_{N,t} + \delta_N^c(\delta), y_{N,t}; \theta_{t+1}^N\right) = u\left(c_{N,t}, y_{N,t}; \theta_t^N\right) + \nu_N \delta \tag{42}$$

and we normalise $\delta_N^y(\delta) = 0$.[31]

By the single-crossing condition higher-type agents see their utility change monotonically through movements along the indifference curve of a lower-type agent, so for small enough δ these equations must solve for unique values of $\delta_n^c(\delta)$ and $\delta_n^y(\delta)$ for all n, with the assumed Inada conditions guaranteeing interiority of the original optimum. These values will preserve incentive compatibility at t. The impact of the perturbations on discounted expected utility from the perspective of prior time periods is left unchanged by the fact that $\sum_{n=1}^N \pi_\Theta(\theta^n)\nu_n = 0$, where this probability measure is common to all true types by the iid assumption.

A.4 Proof of Proposition 5

Our aim is to construct a perturbation schedule $\Delta(\delta)$, to be applied in period t, whose effect effect on the utility of an agent of type θ_t^n will always equal $\nu_n\delta$, and then to consider the marginal impact on the policymaker's resources as δ is moved away from zero. By Lemma 1 we know such a perturbation can be constructed in a manner that preserves incentive compatibility, through the consumption and output perturbations $\delta_n^c(\delta)$ and $\delta_n^y(\delta)$ that are defined in the proof of that Lemma. The net cost of these

[31] This is analogous to the normalisation $\phi^y(\theta, k; c^*, y^*) = 0$ in equation (36).

perturbations on the policymaker's within-period resources in t (per agent with the relevant history) will be:

$$\sum_{n=1}^{N} \pi_{\Theta}\left(\theta_t^n\right)\left[\delta_n^c\left(\delta\right) - \delta_n^y\left(\delta\right)\right]$$

Hence the marginal cost as δ moves away from zero will be:

$$\sum_{n=1}^{N} \pi_{\Theta}\left(\theta_t^n\right)\left[\left.\frac{d\delta_n^c\left(\delta\right)}{d\delta}\right|_{\delta=0} - \left.\frac{d\delta_n^y\left(\delta\right)}{d\delta}\right|_{\delta=0}\right]$$

The rest of the proof is a matter of simple algebra, which we choose to omit, showing that when this object is evaluated and set equal to zero the espression in the Proposition results.[32]

A.5 Proof of Proposition 6

Consider the perturbation given by a movement along the within-period indifference curve of the nth agent, with no changes to the allocations of any other agents. If consumption and output are being jointly reduced this will move us strictly within the constraint set of the relaxed problem, since the net impact on the utility obtainable from reporting the relevant θ^t at t is zero for truth-tellers and strictly negative for 'downwards mimickers' (by single crossing), and expected utility in prior periods is left completely unaffected regardless of the distribution under which it is assessed, by the fact that all agents who remain truth-tellers at t are indifferent to this perturbation. Hence the marginal cost per unit reduction in the utility of potential mimickers must be weakly positive, given that the optimal solution in the relaxed constraint set solves the general problem. From our earlier results, this implies:

$$\frac{\tau\left(\theta_t^n\right)}{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\left(1-\tau\left(\theta_t^n\right)\right)+u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)} \geq 0 \tag{43}$$

[32] Details of the algebra are available on request.

where $u\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)$ (and associated partial derivatives) once more denotes the utility function of an agent whose type is θ_{t+1}^{n+1} mimicking one of type θ_{t+1}^n. We have:

$$\begin{aligned}\left(1-\tau\left(\theta_{t+1}^n\right)\right) &= -\frac{u_y\left(\theta_{t+1}^n\right)}{u_c\left(\theta_{t+1}^n\right)} \\ &> -\frac{u_y\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}{u_c\left(\widehat{\theta}_{t+1}^n;\theta_{t+1}^{n+1}\right)}\end{aligned} \tag{44}$$

where the last inequality is an application of the single-crossing condition. Hence the denominator in condition (43) will be strictly positive, and the result follows.

A.6 Proof of Lemma 2

We need to ensure 'downwards' incentive compatibility continues to hold locally at t and $t+1$. The latter is simpler: it requires that the following conditions are satisfied for agents with the relevant reporting history for all $m \in \{1, ..., N\}$:

$$u\left(c_{m,t+1}^* + \delta_{m,t+1}^c(\delta), y_{m,t+1}^* + \delta_{m,t+1}^y(\delta); \theta_{t+1}^m\right) = u\left(c_{m,t+1}^*, y_{m,t+1}^*; \theta_{t+1}^m\right) + \nu_m\delta \tag{45}$$

$$u\left(c_{m,t+1}^* + \delta_{m,t+1}^c(\delta), y_{m,t+1}^* + \delta_{m,t+1}^y(\delta); \theta_{t+1}^{m+1}\right) = u\left(c_{m,t+1}^*, y_{m,t+1}^*; \theta_{t+1}^{m+1}\right) + \nu_{m+1}\delta \tag{46}$$

where $\delta_{m,t+1}^c(\delta)$ and $\delta_{m,t+1}^y(\delta)$ are the perturbations to the mth agent's consumption and output levels respectively. For the Nth agent we just need:

$$u\left(c_{N,t+1}^* + \delta_{N,t+1}^c(\delta), y_{N,t+1}^*; \theta_{t+1}^N\right) = u\left(c_{N,t+1}^*, y_{N,t+1}^*; \theta_{t+1}^N\right) + \nu_N\delta \tag{47}$$

and we normalise $\delta_{N,t+1}^y(\delta) = 0$.

The proof of Lemma 1 shows that these conditions can indeed be satisfied by appropriate choice of $\delta_{m,t+1}^c(\delta)$ and $\delta_{m,t+1}^y(\delta)$ schedules, given an interior optimum. There remains the problem of incentive compatibility (under the relaxed problem) at t. From the perspective of that time period the $t+1$ perturbations are increasing expected utility for potential mimickers by $\beta\delta$ units, whilst leaving that of truth-tellers constant. To offset this effect we need to move along the indifference curve of the nth agent at t to such an extent that a mimicker's utility is reduced by an offsetting amount. That requires $\delta_{n,t}^c(\delta)$ and $\delta_{n,t}^y(\delta)$ schedules that satisfy:

$$u\left(c_{n,t}^* + \delta_{n,t}^c(\delta), y_{n,t}^* + \delta_{n,t}^y(\delta); \theta_t^n\right) = u\left(c_{n,t}^*, y_{n,t}^*; \theta_t^n\right) \tag{48}$$

$$u\left(c_{n,t}^{*}+\delta_{n,t}^{c}\left(\delta\right),y_{n,t}^{*}+\delta_{n,t}^{y}\left(\delta\right);\theta_{t}^{n+1}\right)=u\left(c_{n,t}^{*},y_{n,t}^{*};\theta_{t}^{n+1}\right)-\beta\delta \tag{49}$$

Again, by the single crossing condition the utility of the agent of type θ_t^{n+1} changes monotonically as one moves along a lower-type agent's indifference curve, so for small enough $|\delta|$ in an open neighbourhood of $\delta=0$ this is always possible.

A.7 Proof of Proposition 7

We consider a composite perturbation pair, denoted $\Delta(\delta)$ and $\Delta_{-1}(\delta)$, such that $\Delta(\delta)$ raises the within-period utility of an agent of type θ_{t+1}^{m} by an amount $\nu_m\delta$ at $t+1$, where ν_m is the mth entry of the vector ν. By earlier arguments (c.f. proof of Proposition 5), the marginal cost of the $\Delta(\delta)$ perturbation as δ is moved away from 0, assessed from the perspective of time t, will be:

$$R_{t+1}^{-1}\left[\sum_{m=1}^{N}\pi_{\Theta}\left(\theta_{t+1}^{m}|\theta_{t}^{n}\right)\nu_{m}MC\left(\theta_{t+1}^{m}\right)-\sum_{m=1}^{N}\pi_{\Theta}\left(\theta_{t+1}^{m}|\theta_{t}^{n}\right)\left(\nu_{m+1}-\nu_{m}\right)DC\left(\theta_{t+1}^{m}\right)\right]$$

This object is equal to (minus) the right-hand side of (25), multiplied by R_{t+1}^{-1}. By Lemma 2 we know that we can remain within the constraint set of the relaxed problem through these perturbations, and the fact that the solution to the relaxed problem also solves the general problem will then imply marginal changes cannot raise a surplus. The proof of Lemma 2 shows that incentive compatibility at t is preserved by moving allocations along the indifference curve of the relevant truth-telling agent with the report history θ^t, and doing so by an amount sufficient to reduce the within-period utility of a mimicker by $\beta\delta$ units. By earlier arguments, the marginal cost of this perturbation as δ is moved away from zero, assessed at time t, will be:

$$\beta\pi_{\Theta}\left(\theta^{t}\right)DC\left(\theta_{t}^{n}\right)$$

The result then follows from the fact that the total present value of the marginal cost of the perturbation must be zero at an optimum.

A.8 Proof of Lemma 4

It remains to establish the result for the case in which consumption and labour supply are Edgeworth complements and productivities are iid. To put a zero lower bound on $MC(\theta_t)$ in this case we need to verify that $\alpha(\theta_t)<1$, and that the denominator of $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}$ is always positive. The latter follows straightforwardly from the definition of $\alpha(\theta_t)$ together with (5). For the former, suppose instead that $\alpha(\theta_t) \geq 1$ were to hold for some θ_t and a given report history. We argue that in this situation it is always possible for the policymaker to generate surplus resources, whilst preserving incentive compatibility.

Suppose we reduce the utility of an agent of type θ^1_{t+1} by a unit *through a reduction in consumption alone*. This will necessarily reduce the utility of an agent of type θ^2_{t+1} who mimicks θ^1_{t+1}, so to preserve exact incentive compatibility at $t+1$ we must reduce the truth-telling utility of θ^2_{t+1} by a compensating amount. Suppose this is likewise done by reducing the consumption of that agent alone. Further reductions in utility must then be provided to θ^3_{t+1}, again assumed to be done through consumption changes alone, and so on up to θ^N_{t+1}. The consumption of all agents will have fallen at $t+1$, and therefore their utility likewise. Suppese that the expected reduction in $t+1$ utility is some amount $\overline{\nu}$. Then incentive compatibility can be preserved from the perspective of period t by raising the within-period utility of the relevant type θ_t by an amount $\beta\overline{\nu}$, whilst ensuring an equal effect on 'one-higher' minickers. But since $\alpha(\theta_t) < 1$ this period-t utility comes at negative cost, whilst by construction the consumption changes at $t+1$ generate positive resources. Hence the combined perturbation generates a surplus, contradicting optimality.

A.9 Proof of Proposition 9

We know Doob's convergence theorem applies to the non-negative martingale $\frac{1-\alpha(\theta_t)}{u_c(\theta_t)+u_y(\theta_t)\alpha(\theta_t)}$, so need only show that it is not possible for this object to converge to any non-zero value. The following Lemma is useful:

Lemma 5 $\frac{\tau(\theta^n_t)}{u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)(1-\tau(\theta^n_t))+u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)} \xrightarrow{a.s.} 0$ *holds under an optimal plan that solves the restricted problem.*

Proof. In the iid case this follows directly from equation (22):

$$\begin{aligned} &\lim_{t\to\infty}\left[-\pi_\Theta\left(\theta^n_{t+1}|\theta_t\right)\frac{\tau\left(\theta^n_{t+1}\right)}{u_c\left(\widehat{\theta}^n_{t+1};\theta^{n+1}_{t+1}\right)\left(1-\tau\left(\theta^n_{t+1}\right)\right)+u_y\left(\widehat{\theta}^n_{t+1};\theta^{n+1}_{t+1}\right)}\right] \qquad (50)\\ &= -\sum_{m=n+1}^{N}\pi_\Theta\left(\theta^m_{t+1}|\theta_t\right)\lim_{t\to\infty}\left[\frac{1-\alpha\left(\theta^m_{t+1}\right)}{u_c\left(\theta^m_{t+1}\right)+u_y\left(\theta^m_{t+1}\right)\alpha\left(\theta^m_{t+1}\right)}\right]\\ &+\pi_\Theta\left(\theta_{t+1}>\theta^n_{t+1}|\theta_t\right)\sum_{\theta_{t+1}\in\Theta}\pi_\Theta\left(\theta_{t+1}|\theta_t\right)\lim_{t\to\infty}\left[\frac{1-\alpha\left(\theta_{t+1}\right)}{u_c\left(\theta_{t+1}\right)+u_y\left(\theta_{t+1}\right)\alpha\left(\theta_{t+1}\right)}\right]\\ &= 0 \end{aligned}$$

In the Markov case we know that equation (22) must hold in periods immediately following those in which $\theta = \theta^N$, and so if one indexes by T the (infinite) set of periods in which this is the case, and denotes by $t(T)$ the (conventional) time period corresponding to the Tth occasion on which $\theta = \theta^N$ has obtained along the given sample path, we must have:

$$\begin{aligned}
&\lim_{T\to\infty} \left[-\pi_\Theta\left(\theta^n_{t(T)+1}|\theta_{t(T)}\right)\right. \\
&\left.\cdot\frac{\tau\left(\theta^n_{t(T)+1}\right)}{u_c\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)\left(1-\tau\left(\theta^n_{t(T)+1}\right)\right)+u_y\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)}\right] \\
&= -\sum_{m=n+1}^{N}\pi_\Theta\left(\theta^m_{t(T)+1}|\theta_{t(T)}\right)\lim_{T\to\infty}\left[\frac{1-\alpha\left(\theta^m_{t(T)+1}\right)}{u_c\left(\theta^m_{t(T)+1}\right)+u_y\left(\theta^m_{t(T)+1}\right)\alpha\left(\theta^m_{t(T)+1}\right)}\right] \\
&+\pi_\Theta\left(\theta_{t(T)+1}>\theta^n_{t(T)+1}|\theta_{t(T)}\right) \\
&\cdot\sum_{\theta_{t(T)+1}\in\Theta}\pi_\Theta\left(\theta_{t(T)+1}|\theta_{t(T)}\right)\lim_{T\to\infty}\left[\frac{1-\alpha\left(\theta_{t(T)+1}\right)}{u_c\left(\theta_{t(T)+1}\right)+u_y\left(\theta_{t(T)+1}\right)\alpha\left(\theta_{t(T)+1}\right)}\right] \\
&= 0
\end{aligned} \tag{51}$$

But if

$$\frac{\tau\left(\theta^n_{t(T)+1}\right)}{u_c\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)\left(1-\tau\left(\theta^n_{t(T)+1}\right)\right)+u_y\left(\widehat{\theta}^n_{t(T)+1};\theta^{n+1}_{t(T)+1}\right)}=0$$

holds at the limit as T becomes large then we must also, at the same limit, have an identical set of zero restrictions in period $t(T)+2$, by equations (25) and (29). By induction this can then be extended to period $t(T)+n$ for all $n > 1$, and the result follows. ■

This Lemma implies two alternatives: either

$$\tau\left(\theta^n_t\right) \overset{a.s.}{\to} 0$$

or

$$u_c\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right)\left(1-\tau\left(\theta^n_t\right)\right)+u_y\left(\widehat{\theta}^n_t;\theta^{n+1}_t\right) \overset{a.s.}{\to} \infty$$

Suppose the latter were true. Expanding out the definition of $\alpha\left(\theta_t^n\right)$ we have:

$$
\begin{aligned}
u_c\left(\theta_t^n\right)+u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right) &= \frac{u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)-u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{u_c\left(\theta_t^n\right)}{u_y\left(\theta_t^n\right)}}{1-\frac{u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)}{u_y\left(\theta_t^n\right)}} \\
&> u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)-u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{u_c\left(\theta_t^n\right)}{u_y\left(\theta_t^n\right)} \\
&= u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)+u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{1}{\left(1-\tau\left(\theta_t^n\right)\right)}
\end{aligned} \tag{52}
$$

If

$$u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\left(1-\tau\left(\theta_t^n\right)\right)+u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\overset{a.s.}{\rightarrow}\infty$$

then

$$u_c\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)+u_y\left(\widehat{\theta}_t^n;\theta_t^{n+1}\right)\frac{1}{\left(1-\tau\left(\theta_t^n\right)\right)}\overset{a.s.}{\rightarrow}\infty$$

must also hold, since $\left(1-\tau\left(\theta_t^n\right)\right)\in[0,1]$ follows from the definition of τ and Proposition 6. Hence we must also have

$$u_c\left(\theta_t^n\right)+u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right)\overset{a.s.}{\rightarrow}\infty$$

This in turn implies $\frac{1-\alpha\left(\theta_t^n\right)}{u_c\left(\theta_t^n\right)+u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right)}$ can only converge to a non-zero limit if $\left|\alpha\left(\theta_t\right)\right|$ is itself always infinite at that limit. But since we know $\alpha\left(\theta_t\right)=0$ when $\theta_t=\theta^N$ we can rule that out.

The alternative is that $\tau\left(\theta_t^n\right)\overset{a.s.}{\rightarrow}0$. In this case we have $u_c\left(\theta_t^n\right)=-u_y\left(\theta_t^n\right)$ at the limit, and so

$$\frac{1-\alpha\left(\theta_t^n\right)}{u_c\left(\theta_t^n\right)+u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right)}=\frac{1}{u_c\left(\theta_t^n\right)}$$

Hence the inverse of the marginal utility of consumption must be converging to a common value for all agents. But since $u_c\left(\theta_t^n\right)=-u_y\left(\theta_t^n\right)$ the marginal disutility of production must also be converging to the *same* value across agents. Suppose this were a finite value. If u_c is common across types and $u_c=-u_y$ holds then it is easy to show that utility must be *decreasing* in type. This is clearly inconsistent with incentive compatibility, which is enough to rule out $\frac{1-\alpha\left(\theta_t^n\right)}{u_c\left(\theta_t^n\right)+u_y\left(\theta_t^n\right)\alpha\left(\theta_t^n\right)}$ converging to a non-zero value in this case too.